STINFO COPY

AFRL-RV-PS-
TR-2007-1196

Magnetic Random Access Memory for Embedded Computing

Gregory W. Donohoe

University of Idaho
Electrical and Computer Engineering
Buchannan Engineering Lab
Moscow, ID 83844-1023

29 October 2007

Final Report

AIR FORCE RESEARCH LABORATORY
Space Vehicles Directorate
3550 Aberdeen Ave SE
AIR FORCE MATERIEL COMMAND
KIRTLAND AIR FORCE BASE, NM 87117-5776

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the Air Force Research Laboratory/RV Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (http://www.dtic.mil).

AFRL-RV-PS-TR-2007-1196 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

//signed//	//signed//
JEFFREY B. SCOTT, 2d Lt USAF Program Manager	JOHN P. BEAUCHEMIN, Lt Col, USAF Deputy Chief, Spacecraft Technology Division Space Vehicles Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
29/10/2007	Final Report	01/05/2004 to 31/07/2007

4. TITLE AND SUBTITLE
Magnetic Random Access Memory for Embedded Computing

5a. CONTRACT NUMBER
FA9453-04-1-0206

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
63401F

6. AUTHOR(S)
Gregory W. Donohoe

5d. PROJECT NUMBER
5159/2181

5e. TASK NUMBER
SC/SC

5f. WORK UNIT NUMBER
AA/AJ

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Idaho
Electrical and Computer Engineering
Buchannan Engineering Lab
Moscow, ID 83844-1023

8. PERFORMING ORGANIZATION REPORT NUMBER
MRAM_II_2007-Final

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Research Laboratory
Space Vehicles Directorate
3550 Aberdeen Ave. SE
Kirtland AFB, NM 87117-5776

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/RVSE

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-RV-PS-TR-2007-1196

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited. (Clearance #VS07-0688)

13. SUPPLEMENTARY NOTES
This report is published in the interest of scientific and technical information exchange. The established procedures for editing reports were not followed for this technical report.

14. ABSTRACT

The goal of this research was to develop an embedded magnetic memory technology to be integrated into Complementary Metal Oxide Semiconductor (CMOS) integrated circuit fabrication process to provide radiation-hard, logic elements and small random-access memories. The goal is not to provide large-scale, bulk memory, but latches and flip flops that serve as state and data registers for sequential logic, and configuration registers for configurable logic. The benefits to spacecraft systems include the ability to power-down a subsystem while retaining system state, thus saving energy until the subsystem is required. The subsystem can then be powered-up and begin operating in milliseconds. The technology is based on a unique, PacMan-shaped magnetic tunneling junction (MTJ) cell developed at the University of Idaho. The focus of this research is to refine the PacMan cell to make it practical for integration into CMOS circuits, to develop CMOS circuits that employ the magnetic cells, and to integrate the cells onto a CMOS process. The process produced two circuit designs based on magnetic memory elements: a magnetic latch, and a magnetic shadow memory to serve as a backup to volatile electronic memory.

15. SUBJECT TERMS
CMOS, MTJ Cell

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Jeffrey B. Scott
Unclassified	Unclassified	Unclassified	Unlimited	212	**19b. TELEPHONE NUMBER** *(include area code)* (505) 846-6280

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. 239.18

Magnetic Random Access Memory for Embedded Computing

From: Gregory W. Donohoe, PhD, PE
Electrical and Computer Engineering
University of Idaho
Moscow, ID 83844-1023

To: Lt. Jeffery Scott
AFRL/RVSE
Kirtland AFB, NM 87117-5776
Jeffery.Scott@kirtland.af.mil

Re: Final Technical Report for contracts FA9453-04-1-0206-P000001 and FA9453-04-1-0206-P000002

Period of Performance: 1 May 2005 to 31 July 2007

Report Number: MRAM_II_2007-Final

Table of Contents

Abstract 2

List of Figures 2

1. Introduction 3

1.1 Innovation 3

1.2 Benefits 3

1.3 Approach 4

1.4 Magnetic Tunneling Junction Memory Cells 4

1.4.1 University of Idaho Magnetic Element 4

1.4.2 Basic MTJ Operation 4

1.4.3 New Element Shape and Magnetic States 5

2. MTJ Cell Development 6

2.1 Optimized PacMan “soft” magnetic element 6

2.2 Complete MTJ Cells 7

3. Storage Circuits Employing Magnetic Memory Cells 8

3.1 One-wire Writing Scheme 9

3.2 Single-wire Programming in Process Integration 10
3.2 Single-Wire Simulations 11
4. Integration into CMOS circuits 11
4.1 Results of Integration Study 14
5. Conclusions 16
6. References 17
Appendix 1: List of Symbols, Abbreviations and Acronyms 19
Attachment 1: Doctoral Dissertation 19

Abstract

The goal of this research was to develop an embedded magnetic memory technology to be integrated into a Complementary Metal Oxide Semiconductor (CMOS) integrated circuit fabrication process to provide radiation-hard, logic elements and small random-access memories. The goal is not to provide large-scale, bulk memory, but latches and flip flops that serve as state and data registers for sequential logic, and configuration registers for configurable logic. The benefits to spacecraft systems include the ability to power-down a subsystem while retaining system state, thus saving energy until the subsystem is required. The subsystem can then be powered-up and begin operating in milliseconds. The technology is based on a unique, PacMan-shaped magnetic tunneling junction (MTJ) cell developed at the University of Idaho. The focus of this research is to refine the PacMan cell to make it practical for integration into CMOS circuits, to develop CMOS circuits that employ the magnetic cells, and to integrate the cells onto a CMOS process. The produce produced two circuit designs based on magnetic memory elements: a magnetic latch, and a magnetic shadow memory to serve as a backup to volatile electronic memory.

List of Figures

Figure 1: Magnetic Tunneling Junction stack 4
Figure 2: Various types of "Pac Man" shaped elements. 5
Figure 3: Micromagnetic simulation showing single domain reversal (NiFe, 1.05 um x 0.25 um x 40 nm) . 6
Figure 4: Final PacMan soft magnetic element. 6
Figure 5: MTJ cell fabricated between orthogonal electrodes 7
Figure 6: Test array of MTJ cells 7
Figure 7: Differential magnetic flip flop. 8

Figure 8: Magnetic shadow RAM flip flop. 9
Figure 9: Orientation of MTJ cells for opposite polarity. 9
Figure 10: One-wire writing configurations. 9
Figure 11: Single-wire MTJ cell arrangment 10
Figure 12: X-Y MTJ cell arrangement 10
Figure 13: Proposed MTJ cell structure for single-wire programming (not to scale). 11
Figure 14: Test structure for process integration. 12
Figure 15: Short loop test structure. 12
Figure 16: Process for fabricating MTJ cells on short-loop structures. 13
Figure 17: Test structures fabricated on American Semiconductor metal. 14
Figure 18: Plot of surface roughness of FlexFET 6K metal 15

1. Introduction

The purpose of this work was to develop an integrated magnetic memory technology, to provide low-power, nonvolatile, radiation-hard memory on a complementary metal oxide semiconductor (CMOS) process. This provides the means to create nonvolatile latches, flip flops, small amounts of random access memory, data registers, and nonvolatile lookup tables.

1.1 Innovation

This research is based on three innovative elements: (1) the University of Idaho's unique magnetic element shape, which promises improved repeatability and reduced power consumption for memory writes; (2) integration of magnetic memory components into a commercial integrated circuit design flow; and (3) an analysis of the uses and benefits of embedded magnetic memory for reconfigurable architectures.

1.2 Benefits

Radiation-hard, nonvolatile memory used in strategic parts of electronic systems offer increased responsiveness and reduced power consumption. A processor that uses nonvolatile memory for primary off-chip storage does not need to be "booted" after it is powered down; it can be powered back up in an "instant on" state, saving startup time and power. Nonvolatile Magnetic Tunneling Junction (MTJ) memory can provide this "instant on" capability. The magnetic storage cells themselves consume no power when not being accessed, and are inherently radiation hard. MTJ has the potential to serve as a key enabling technology for low-power processing and intelligent power management aboard power-critical systems, such as spacecraft. This will increase the quantity and quality of on-board data processing for a fixed power budget, enabling operational spacecraft to gather more and higher-quality data.

1.3 Approach

The major elements of the project are:

1. Develop and optimize a functional magnetic storage (MTJ) cell
    a. Develop an optimized "free" layer, whose magnetic orientation can be easily switched by a magnetic field.
    b. Develop a complete MTJ cell, incorporating the optimized free layer, whose resistance can be established by switching the magnetic polarity of the free layer.
2. Develop and model data storage circuits based on the MTJ cells.
3. Integrated the MTJ cells into a CMOS process.

The CMOS process chosen is the FlexFET™ Silicon-on-Insulator (SOI) developed by American Semiconductor, Inc., of Boise, ID, under the auspices of the Air Force Research Laboratory.

1.4 Magnetic Tunneling Junction Memory Cells

1.4.1 University of Idaho Magnetic Element

The University of Idaho has patented a unique shape of magnetic element for MTJ cells that shows superior qualities when compared to other shapes for MTJ structures. Dubbed the PacMan, key performance metrics are high selectivity and narrow switching field distribution. Narrow switching field distribution leads to high repeatability and high yield: we can design circuits to generate the necessary switching fields predictably and reliably.

1.4.2 Basic MTJ Operation

The basic MTJ cell in MRAM device is composed of four successive layers: a "pinning" layer (antiferromagnetic material), a "pinned" layer of fixed magnetic orientation (ferromagnetic material), an ultrathin barrier layer (insulator), and a "free" layer (ferromagnetic material), as shown in Figure 1. The magnetization state of the free layer can be placed in one of two states, parallel or antiparallel to the magnetization state of the pinned layer. If the magnetic fields of both the pinned and free magnetic layers are aligned, then the resistance from the top to the bottom of the stack is reduced: spin-polarized tunneling through the insulating barrier layer. By placing the free layer in a magnetic state opposite to that of the pinned layer, the resistance through the stack will br increased. The difference in high and low resistance, or ΔR,

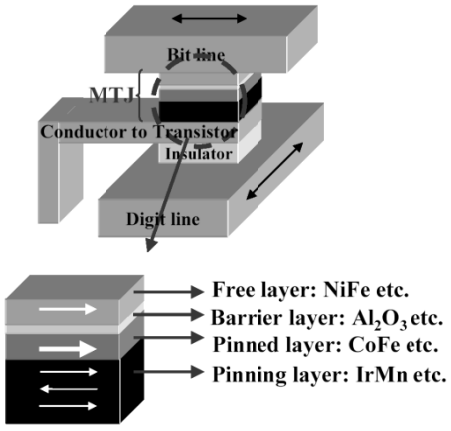

Figure 1: Magnetic Tunneling Junction stack

can range from 5 to 50%. Readout circuitry detects this difference for a logical 0 or 1.

Desirable qualities for an MTJ cell are:

- For writing properties
  - High selectivity of a cell addressed by two orthogonal magnetic fields
  - Switching repeatability of the free layer, i.e. the ability to reliably place the free layer in the desired state with a predictable writing current;
  - Low switching field distribution of free layer between MTJ elements
  - Low power consumption for writing
  - High switching speed
  -
- For reading properties
  - High reading speed
  - High ΔR
  - Uniformity of junction resistance (small variance) between MTJ elements on a wafer, which enables reliable readout
  - Thermal stability.
  - Large number of switching cycles before the mechanism wears out

1.4.3 New Element Shape and Magnetic States

A current issue in magnetic memories is how to achieve a perfect selectivity of a MTJ cell by two orthogonal magnetic fields. In order to address this issue, there are significant challenges to overcome: a wide switching field distribution due to variation of end shape between elements, interlayer magnetic interaction, and unrepeatable switching of the free layer in a MTJ cell caused by various magnetic defects, which involve during switching, such as edge domains, domain walls, or vortices.

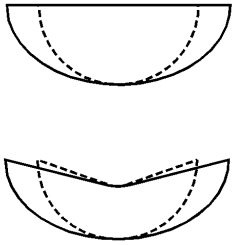

Figure 2: Various types of "Pac Man" shaped elements.

The University of Idaho MT cell based on a unique material composition and shape that promises superior performance in MTJ cells, so-called “Pac Man” shape (Fig 2).

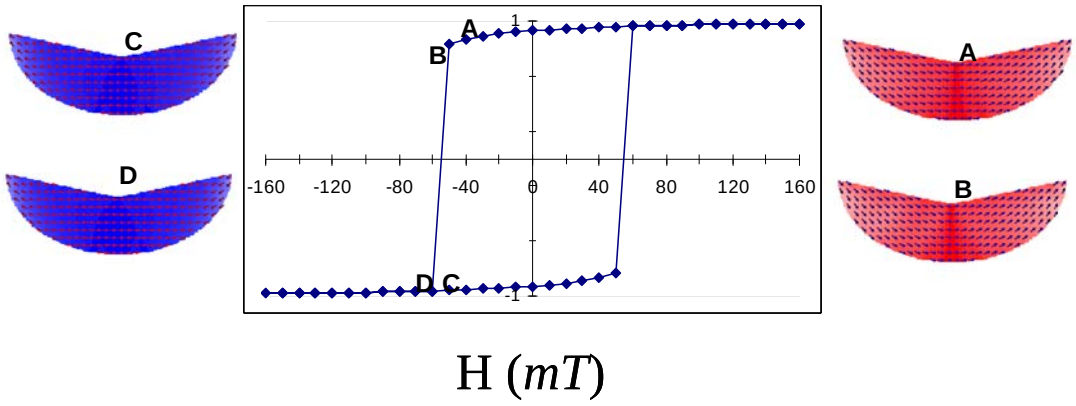

Figure 3: Micromagnetic simulation showing single domain reversal (NiFe, 1.05 um x 0.25 um x 40 nm)

The magnetic orientation is reversible in a repeatable way and has a lower switching field distribution in an array as compared to the hexagon elements with the overall same dimensions (Fig. 3). In addition, its magnetization configuration resembles C state, which has been known to have a high selectivity. Therefore, it is expected that the switching of the free layer will occur at a low easy axis magnetic field with an assistance of a comparatively small hard axis magnetic field (refer to *Appl. Phys. Lett.* ***83****, 329 (2003)*).

As the size of magnetic element for a MTJ scales, the superparamagnetic phase emerges in which the room temperature thermal energy overcomes the magnetostatic energy of the element, resulting in zero magnetic hysteresis. In other words, although the magnetic element is a single-domain ferromagnet, the ability of ferromagnetic film to retain magnetization is lost when its dimension is below a threshold. However, the thermal stability is achieved by designing layered film structures and composition of ferromagnetic materials. In theory, there is no limit to the number of times MTJ cells can switch; there is no known wear-out mechanism.

For this research, a fifth property is required: radiation hardness. The MTJ cells themselves are intrinsically radiation hard. We will need to develop radiation hard read and write electronics to produce a fully radiation-hard system.

2. MTJ Cell Development

2.1 Optimized PacMan "soft" magnetic element

The first challenge was to optimize the shape and material composition of the "free" magnetic layer to enable fast, reliable switching with a small amount of energy. The unique properties of the PacMan shape enable it to be programmed (i.e., force into a particular magnetic polar orientation) repeatedly.

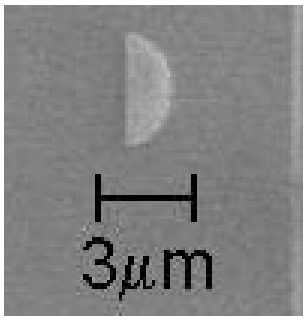

Figure 4: Final PacMan soft magnetic element.

Altering the details of the shape changes its switching properties, as does altering the material composition and size.

The team experimented with a variety of element compositions and shapes to achieve a set that had the desired properties. The primary parameter adjusted was the angle of the PacMan "mouth" opening. The optimal shape turned out to be defined by an arc on one side and a straight line on other, as shown in the figure.

2.2 Complete MTJ Cells

Arrays of PacMan cells of different sizes and material compositions were designed and fabricated. Figures 5 and 6 show a representative examples. The cells were fabricated between two orthogonal copper electrodes so their electrical properties could be read in a probe station. The cells were programmed with bulk magnetic fields (not electrically programs) and their properties measured. The resulting cells had tunneling magnetic resistance ratios (TMR) ranging from 15% to 35%, and a nominal resistance of 10 to 15 Kilohms.

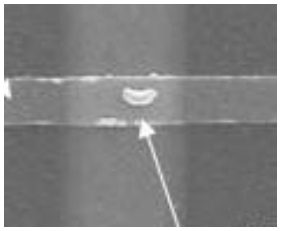

Figure 5: MTJ cell fabricated between orthogonal electrodes

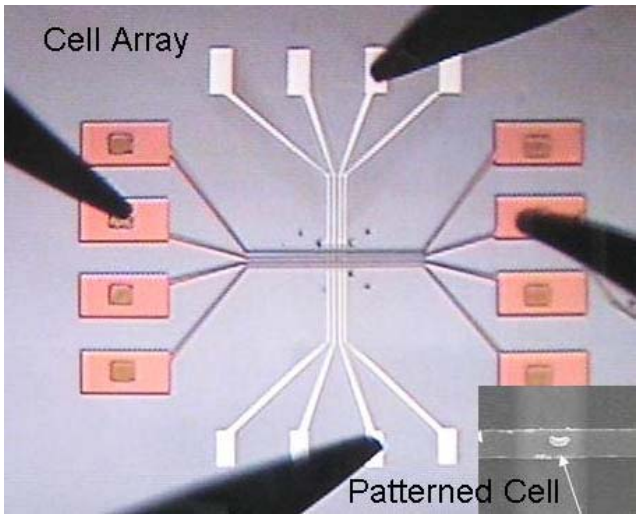

Figure 6: Test array of MTJ cells

The cells had write currents ranging from 10 to 25 mA in short pulses, and switching times on the order of one to two nanoseconds. The write currents are higher than we would like, and we are working on ways to reduce them.

3. Storage Circuits Employing Magnetic Memory Cells

Two kinds of circuits were designed and simulated.

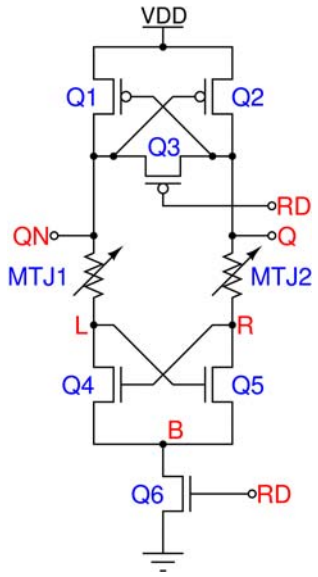

Figure 7: Differential magnetic flip flop.

- A differential magnetic flip flop
- A magnetic shadow flip flop

A circuit diagram of the differential magnetic flip flop is shown in Figure 7. MTJs are placed between the P and N halves of a 4-transistor latch The MTJs are programmed to opposite states In case of a single event upset, the flip flop will restore the original state as follows: One of the signals, L or R will fall faster than the other, due to the different in MTJ resistance. Positive feedback will force the appropriate output, Q or QN, to be low. The MTJs provide resistor isolation for SEU immunity

The magnetic flip flop takes advantage of the back gate of the FlexFET process. The circuit is shown in Figure 8. The MTJs are programmed to opposite states. During a restore operation, the back gates discharge at different rates, causing different voltages in Q1 and Q2. The MTJs operate in common mode.

This scheme is easily applied to a Single-Event Upset (SEU) immune cell such as the widely available DICE cell.

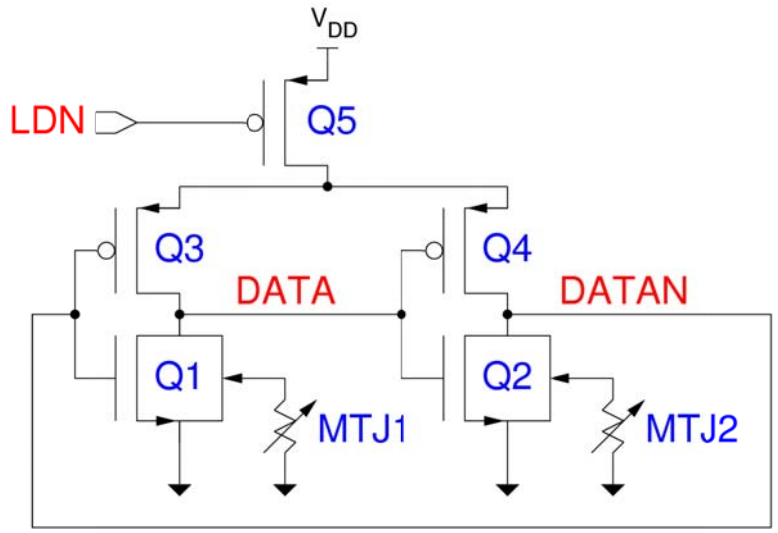

Figure 8: Magnetic shadow RAM flip flop.

3.1 One-wire Writing Scheme

Both of these circuits use one-wire writing. In this scheme, we abandon the X-Y addressing commonly employed in MRAMs. For the flip flop, we want the two MTJ cells to be programmed with opposite magnetic polarities. We achieve this by arranging the cells as shown in Figure 9. The write line is folded so the MTJ cells will experience magnetic fields of opposite senses from the same current.

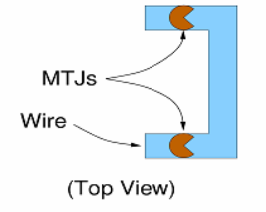

Figure 9: Orientation of MTJ cells for opposite polarity.

Figure 10(a) shows the magnetic fields (in green) from the electrical configuration shown, where current moves down the vertical metal run. In Figure 10(b), the writing current is reversed, and so are the magnetic fields generated by the currents. The switching scheme in Figure 10(c) can be implemented using drive transistors. In cases where banks of MTJ cells are to be

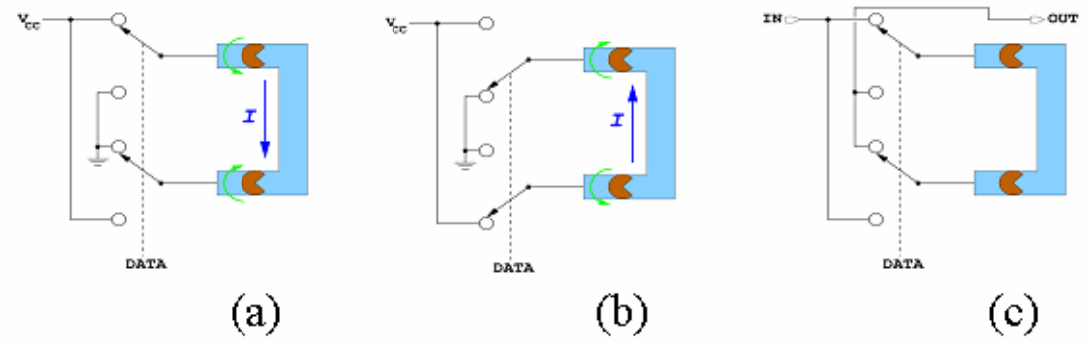

Figure 10: One-wire writing configurations.

written, as in data registers, the flip flops can be programmed in series, thus sharing a write current and minimizing the power required.

The relatively small TMR of the MTJ cells and inevitable process variations suggested a differential circuit structure. We designed and simulated several versions, in the form of magnetic flop flops and shadow RAMs An interesting observation is that pairs of magnetic elements positioned in certain arrangements can share magnetic fields either constructively or destructively, which might be exploited to reduce write current [17].

3.2 Single-wire Programming in Process Integration

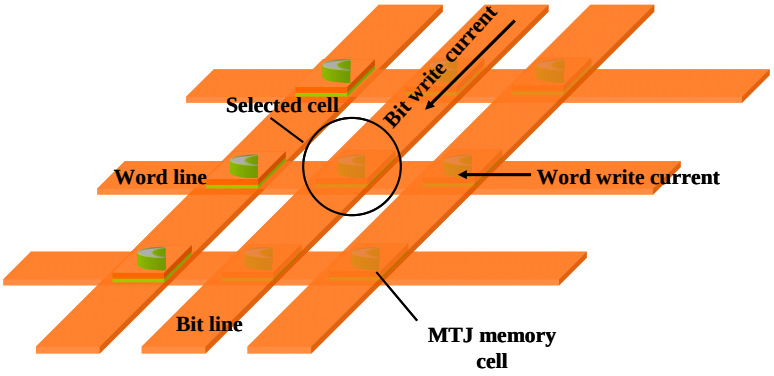

Figure 12: X-Y MTJ cell arrangement

The usual MRAM architecture uses *two-wire programming*. It employs an X-Y array of MTJ cells, with a parallel set of metal conductors above the cells (e.g. bit lines) and an orthogonal set of conductors beneath (word lines), as shown in Figure 1. The MTJ cells a are arranged with their magnetic axes at a forty-five-degree angle to both lines. A cell is selectively written by driving its word and bit lines with a current that generates the electromagnetic fields whose vector sum causes the magnetic polarization of the selected cell to switch. It is important that no other cell in the matrix is induced two switch – that is, the cells must be highly *selective*.

The two-wire programming scheme presents a challenge for back-end CMOS processing. It is desirable to use existing metal layers to program the cells. However, at reasonable currents, on the order of 20 mA, the required distance between the conductors and the cells is small, on the order of ten microns. The MTJ cell is about 40 nm tall. The space between metal layers in the 0.25 micron CMOS process is about 100 microns. The cells cannot be placed between metal layers such that they are close enough to both metal lines to generate a strong enough electromagnetic field at reasonable currents.

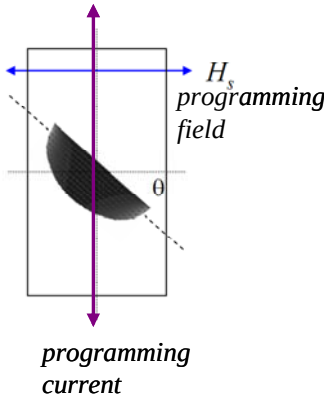

Figure 11: Single-wire MTJ cell arrangment

Two possible solutions are to develop a back-end metal process that meets the space requirements, or to abandon the X-Y two-wire scheme, and use single-wire programming. This project pursued the latter option.

3.2 Single-Wire Simulations

Figure 11 shows the arrangement for single-wire programming. For optimal programming, the MTJ cell must be placed at the proper angle with respect to the magnetic field, and hence to the programming, which is orthogonal to the field. The optimal angle is generally accepted to be about 45 degrees. To investigate ideal placement and assess feasibility, we modeled an MTJ cell at different angles with respect to the current.

The Pac Man I (PM I) shape was chosen, in which the notch is opened to 180 degrees, or a straight line. The experimental parameters are:

Proposed Scenario

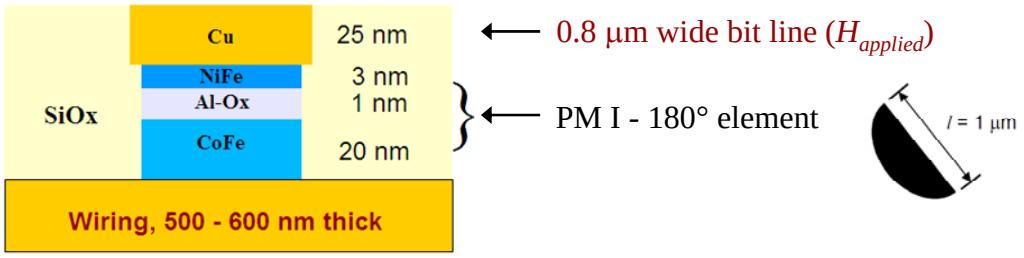

Figure 13: Proposed MTJ cell structure for single-wire programming (not to scale).

A workable scenario is to $1\,\mu m \times 0.36\,\mu m$ Pac Man I 180-degree cell with the cross-section shown, at 30 degrees orientation, sandwiched between an aluminum conductor and a copper write line. With this arrangement, a current pulse of 11.8 mA, lasting about 2.5 ns, creates a magnetic field strength of 90 Oersteds, sufficient to completely switch the cell. For logic purposes, the switching is virtually complete after 1.2 ns.

4. Integration into CMOS circuits

American Semiconductor, Inc., of Boise was selected as a CMOS integrated circuit partner. Their radiation-hard FlexFET™ Silicon-on-Insulator CMOS process was developed under the guidance of AFRL/VSSE for aerospace applications. Unlike large IC production companies, American Semi specializes in custom, developmental work of this sort, and proved to be very willing to work with us.

The FlexFET™ process has the following advantages:

- High-performance "Silicon on Insulator" technology
- Ability to adjust operating voltage, which is exploited in another University of Idaho program for ultra-low-power microelectronics. This is the technology that can benefit most from our integrated passive components.
- Radiation hard for space applications.
- Fabricated at an on-shore U.S. facility (Cypresss Semiconductor in California)

The first step was to determine whether MTJ cells can actually be fabricated on on American Semiconductor's metal.

To determine this, a test array was designed and fabricated in a "short loop" cycle on the American Semiconductor process. The short loop cycle produces test wafers, but does not go through the entire integrated circuit fabrication process. In this case, the short loop tests only go through the field oxide and metal deposition steps. Figure 14 below shows the test structure. It consist of horizontal and vertical traces. The horizontal traces represent the lower metal level, and the vertical traces the upper level. The vertical traces are not complete: they leave a gap where the MTJ cell array is to be fabricated.

After the metal structures are fabricated, the wafers are removed from the fabrication line and diced. In the lab, the MTJ cells are deposited, and the upper metal traces completed.

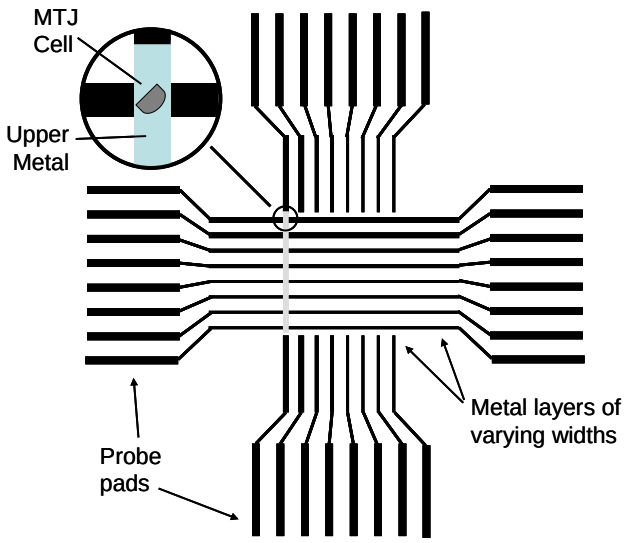

Figure 14: Test structure for process integration.

Figure 15 shows the process used to fabricate the test cells onto the FlexFET metal. Figure 16 is a photograph of the metal test structures. Figure 17 is an Atomic Force Microscope image of the MTJ cells fabricated on the metal test structures. Figure 18 is a plot of the magnetic tunneling resistance, showing very low resistance (around 26 ohms), and very small resistance ratio (less than 1%),

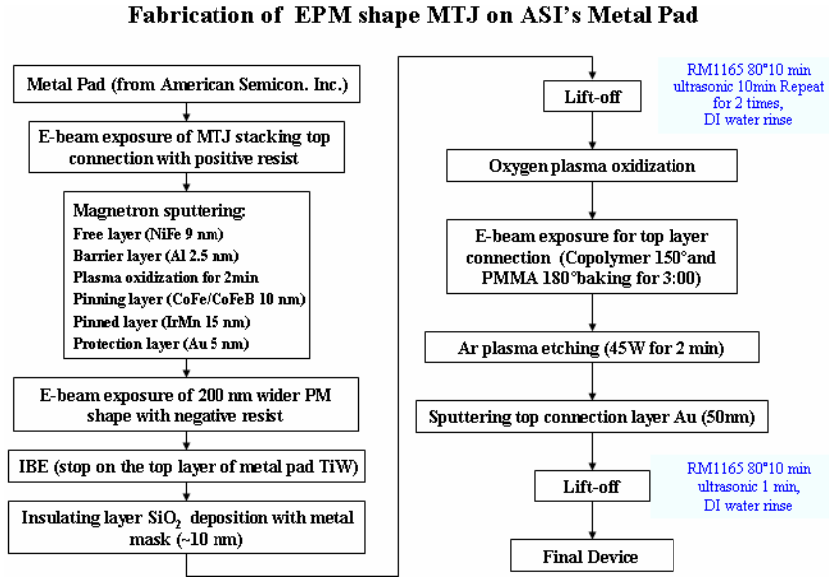

Figure 16: Process for fabricating MTJ cells on short-loop structures.

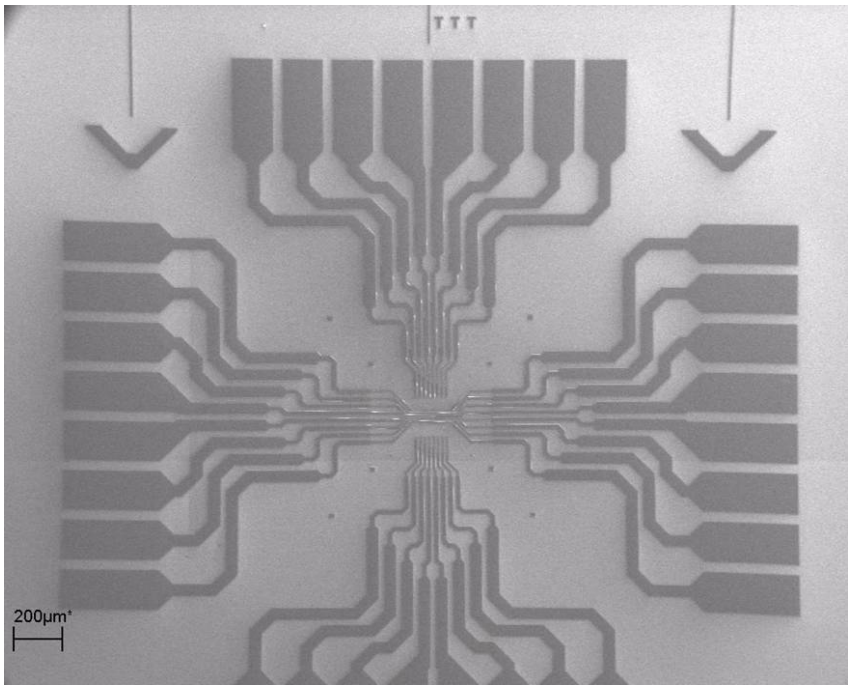

Figure 17: Test structures fabricated on American Semiconductor metal.

4.1 Results of Integration Study

MTJ cells were successfully fabricated on the FlexFET metal, but the performance of these cells was disappointing.. The first issue to address is the surface roughness of the FlexFET metal, which seems to affect MTJ cell performance. The MTJ cells are only about 40 nm high, comparable to the roughness of the metal. They MTJ cells tend to follow the contour of the metal, and the resulting shape destroys their desired electrical properties.

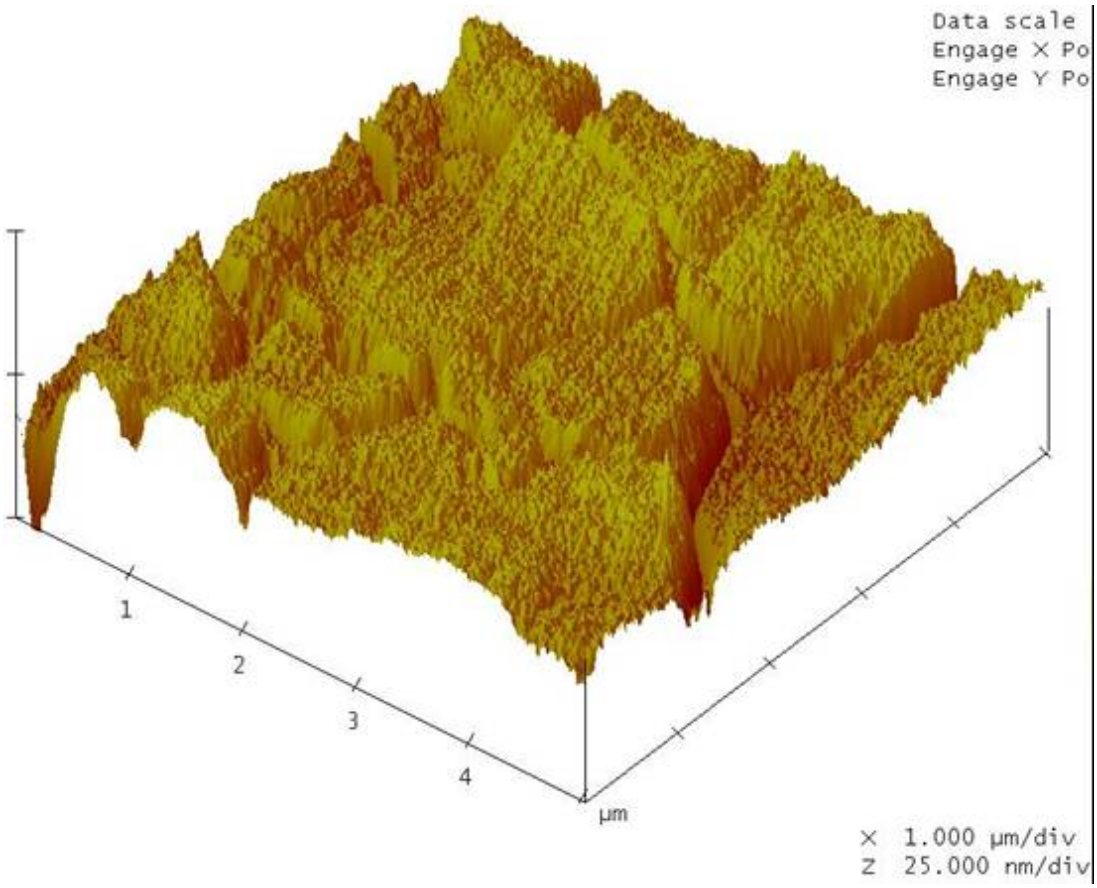

Figure 18: Plot of surface roughness of FlexFET 6K metal

Figure 19 is a cross-sectional display and analysis of the metal. The grain boundaries of the underlying aluminum are clearly visible.

The team considered the follow choices to mitigate the roughness problem:

1. Select a different metal layer. The top metal (20K) is smoother than the others, so we should work with it first.
2. Polish the surface with Chemical-Mechanical Polishing (CMP)
3. "Decorate" the grain boundaries and etch back. For example, we have used spin-on glass (SOG), followed by ion beam etching.
4. Optimization of the Al deposition temperature.
5. Select a different material on which to fabricate MTJ cells.

Option #2 has showed little success. Option 3 has reduced the RMS roughness from 4.65 to 3.16 nm, and cut the peak in half from 28.88 to 15.23 nm. Option #4 and #5 require a new fabrication runs. The team decided to pursue option #5, fabricating MTJ cells onto a tungsten surface, which is much smoother than aluminum. This was carried out under a separate contract, and the results will be reported accordingly. The results were encouraging.

5. Conclusions

The project produced workable magnetic tunneling junction (MTJ) memory cells, which were programmed in a bulk magnetic field. In the process, the researchers learned and published many details about the switching behavior of magnetic devices, which have an impact on the optimal design of these devices. A feasible set of latch and shadow RAM circuits was developed and modeled.

The specifications for the MTJ cells required for successful implementation of the magnetic latches is:

- Nominal resistance: 10 KΩ
- Tunneling Magnetic Resistance Ratio (TMR): 15%
- The preferred size and shape is an elongated PacMan, 1 µm in length

Integrating the MTJ cells onto a CMOS process proved elusive because of the difficulties in creating a metal surface smooth enough to accept the cells. This is necessary in order to carry out the next phases of the research, i.e., to develop accurate electronic circuit (SPICE) models, to implement the on-wire writing scheme, and finally, to integrate MTJ cells with CMOS electronics. This work is being carried out in a follow-on project

Since patterned test cells were not available, due to the roughness problem described above, it has not been possible to test the one-wire writing concept. Simulations indicate that the write current pulses will have to be 20 mA for about 3 ns. This is much higher than desired. Some approaches will be studied to enable writing the cells with much lower current. The most promising is to clad the cells with a magnetic material that focuses the magnetic field onto the cell, where it is needed.

6. References

1. H Xiao, Q.F., B.C. Choi, J. Rudge, Y.K. Hong and G.W. Donohoe, "Vortex core dynamics in submicron ferromagnetic disks", to appear, *Proc. Joint Magnetism and Materials/Intermag Conference (AIP)*, Baltimore, MD, Jan 7-12, 2007.

2. K.J. Hass, G.W. Donohoe and Y.K Hong, "Integrated Magnetic Memory for Embedded Computing Systems, accepted for the *IEEE Aerospace Conference*, Big Sky, MT, March 3-10, 2007.

3. Hass, K.J, G.W. Donohoe, Y.-K. Hong, B.-C. Choi, and K. DeGregorio, "Magnetic shadow RAM", to appear, *Proc. 7th Annual Nonvolatile Memory Symposium*, Nov. 5-8, San Mateo, CA, 2006.

4. K.J. Hass, G.W. Donohoe, Y.-K. Hong and B.C. Choi, "Magnetic flip flop for space applications", IEEE Trans Magnetics, Vol. 42, No. 10, pp. 2751-2753, October 2006.

5. Y.K. Hong, M.H. Park, B.C. Choi, H. Han*, S.H. Gee*, and G.W. Donohoe, "Thickness dependence of vortex head-to-head domain walls in Ni80Fe20 ring elements", in review, Journal of Magnetism and Magnetic Materials, July 2006.

6. Q.F. Xiao*, J. Rudg*e, B.C. Choi, Y.K. Hong and G.W. Donohoe, "Dynamics of Vortex Core Switching in Ferrmagnetic Nanodisks", in review, Applied Physics Letters, July 2006.

7. Q.F. Xiao*, J. Rudge*, B.C. Choi, Y.K. Hong, and G.W. Donohoe, "Dynamics of ultrafast magnetization reversal in submicron elliptical Permalloy thin film elements", Physical Review B 73, 1 (2006).

8. J. Jabal*, Y. K. Hong, B. C. Choi, H. Han, S. H. Gee, G. Abo, K. J. Hass, and G. W. Donohoe, "*Lateral size dependence of Pac-man shaped $Ni_{80}Fe_{20}$ element on magnetization reversal*," accepted for presentation at IEEE International Conference Magnetics Conference, San Diego, CA., May 8 - 12, 2006

9. B. C. Choi, Q. F. Xiao, Y. K. Hong, S. H. Gee*, J. Jabal*, H. Han, K. J. Hass, and G. W. Donohoe, "*Numerical simulation study of magnetization precession dynamics in submicron elliptical $Ni_{80}Fe_{20}$ thin film elements*," accepted for presentation at IEEE International Conference Magnetics Conference, San Diego, CA., May 8 - 12, 2006

10. Y.K. Hong, M.H. Park, B.C. Choi, S.H. Gee and G.W. Donohoe, "Thickness Dependence of Vortex Head-to-Head Domain Walls in $Ni_{80}Fe_{20}$ Ring Elements", submitted to the International Conference on Magnetism, Kyoto, Japan, Aug 21-25, 2006.

11. K.J. Hass, G.W. Donohoe, Y.K. Hong, B.C. Choi, "Magnetic Latches for Space Applications", accepted, IEEE International Magnetics Conference, 2006.

12. J. Jabal*, Y.K. Hong, B.C. Choi, H. Han*, G. Abo*, S.H. Gee* and G.W. Donohoe, “Suppression magnetization ringing in submicron Pac-man shaped Ni80Fe20 thin film elements”, in review, IEEE Conference on Electromagnetic Field Computation”, January 2006.

13. J. Jabal, et al, “Lateral Size Dependence of Pac-man shaped $Ni_{80}Fe_{20}$ Element on Magnetization Reversal”, accepted, IEEE International Magnetics Conference, 2006.

14. K.J. Hass, G.W. Donohoe, Y.K. Hong, B.C. Choi, S. Parke and K. DeGregorio, “Magnetic Shadow RAM”, in review, IEEE VLSI Circuits Symposium, January 2006.

15. B.C. Choi, J. Ho, Y.K. Hong, M.H. Park, H. Han*, S.H. Gee* and G.W. Donohoe, “Magnetization Switching Dynamics Depending on As-Patterned Magnetization State in Magnetic Thin Film Elements”, *IEEE Trans. Magnetics*, Vol. 41, No. 10, October, 2005, pp. 2709-2711.

16. B.C. Choi, B.R. Pujada, Y.K. Hong, M.H. Park, H. Han*, S.H. Gee*, G.W. Donohoe, “Mircromagnetic Domain Structures and Magnetization Switching Mechanism in Submicron Thin Film Elements”, *IEEE Trans. Magnetics*, Vol. 41, No. 10, October, 2005, pp.3109-3111.

17. H. Han, Y.K. Hong, M.H. Park, B.C. Choi, S.H. Gee*, J.F. Jabal, G. Abo*, A. Lyle*, B. Wong* and G.W. Donohoe, “Interaction Effect on Switching Behaviors of Paired “Pac-man” Array, in review, *International Magnetics Conference,* Nagoya, Japan, April 4-8, 2005.

18. B.C. Choi, Y.K. Hong, M.H. Park, H.Han*, S.H. Gee* and G.W. Donohoe, “Magnetization Switching Dynamics Depending on As-Patterned Magnetization States in Magnetic Thin Film Elements”, *International Magnetics Conference*, Nagoya, Japan, April 4-8, 2005.

19. B.C. Choi B.R. Pajada, Y.K. Hong, M.H. Park, H. Han*, S.H. Gee*, G.W. Donohoe, “Micromagnetic Domain Structures and Magnetization Switching Mechanism in Submicron Thin Film Elements”, *International Magnetics Conference*, Nagoya, Japan, April 4-8, 2005.

20. B.P. Puhada, B.C. Choi, M.H. Park, Y.K. Hong, S.H. Gee*, H. Han*, and G.W. Donohoe, “Micromagnetic configurations and switching mechanism in Pac-man shaped $Ni_{80}Fe_{20}$ submicron magnetics”, *Journal of Applied Physics*, 97, 073904 (2005).

21. B. C. Choi, Q. F. Xiao, Y. K. Hong, S. H. Gee*, J. Jabal*, H. Han*, K. J. Hass*, and G. W. Donohoe, “Numerical simulation study of magnetization precession dynamics in submicron elliptical $Ni_{80}Fe_{20}$ thin film elements,” *IEEE International Magnetics Conference,* San Diego, CA., May 8 - 12, 2006

22. K.J. Hass, G.W. Donohoe, Y.K. Hong, B.C. Choi, “Magnetic Latches for Space Applications”, *IEEE International Magnetics Conference*, May, 2006.

Appendix 1: List of Symbols, Abbreviations and Acronyms

CMOS	Complementary Metal-Oxide Semiconductor
CMP	Chemical-mechanical polishing
DICE cell	Dual-Interlocked Cell – a single-even-upset resistant latch
FlexFET™	American Semiconductor "Flexible Field Effect Transistor" process
MRAM	Magnetic Random Access Memory
MTJ	Magnetic Tunneling Junction
RAM	Random Access Memory
SEU	Single Event Upset (Radiation effect)
SOG	Spin-on Glass
SPICE	Simulation Program with Integrated Circuit Emphasis
SOI	Silicon-on-Insulator, a form of CMOS structure
TMR	Tunneling Magnetic Resistance Ratio

Attachment 1: Doctoral Dissertation

Kenneth Joseph Hass, Radiation-Tolerant Embedded Memory using Magnetic Tunneling Junctions, PhD Dissertation, University of Idaho, 2007.

RADIATION-TOLERANT EMBEDDED MEMORY USING MAGNETIC TUNNEL JUNCTIONS

A Dissertation

Presented in Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

with a

Major in Electrical Engineering

in the

College of Engineering

University of Idaho

by

Kenneth Joseph Hass

May 2007

Major Professor: Gregory W. Donohoe, Ph.D.

AUTHORIZATION TO SUBMIT DISSERTATION

This dissertation of Kenneth Joseph Hass, submitted for the degree of Doctor of Philosophy with a major in Electrical Engineering and title "Radiation-Tolerant Embedded Memory Using Magnetic Tunnel Junctions," has been reviewed in final form. Permission, as indicated by the signatures and dates given below, is now granted to submit final copies to the College of Graduate Studies for approval.

Major Professor _____ Date _____
Gregory Donohoe

Committee Members _____ Date _____
James Frenzel

_____ Date _____
Yang-Ki Hong

_____ Date _____
Kenneth Noren

Department Administrator _____ Date _____
Brian Johnson

Discipline's College Dean _____ Date _____
Aicha Elshabini

Final Approval and Acceptance by the College of Graduate Studies

_____ Date _____
Margrit von Braun

Abstract

Non-volatile memories occupy an important niche in the universe of solid state memory devices. They are able to retain their programmed state indefinitely without power, but can be reprogrammed as desired *in situ*. Aerospace applications have a significant need for non-volatile memories but they also present requirements that are not encountered in commercial electronics. In particular, many aerospace systems must be able to operate continuously and flawlessly in the natural radiation environment of space.

The recent development of practical spin-polarized magnetic tunneling junctions has made possible a new class of non-volatile memories, known collectively as magnetic RAM (MRAM). Information is stored as the orientation of magnetic fields in sub-micron ferromagnetic elements, which is expected to provide much higher resistance to the effects of ionizing radiation than memory technologies that rely on stored charge.

While commercial semiconductor designers are actively pursuing the development of bulk MRAM, where millions of bits of memory are incorporated into a single integrated circuit, there has been little research devoted to integrating magnetic memory elements with logic circuits. The goal of this research is to design such embedded magnetic memories with the additional requirement that they must be highly resistant to ionizing radiation.

This dissertation begins with an overview of the quantum phenomena at work in magnetic tunneling junctions, the natural radiation environment of space, and radiation effects in CMOS electronics. This leads into a discussion of the goals, constraints, and obstacles that must be considered when designing radiation-tolerant embedded magnetic memories. Prior art in this area is presented and evaluated.

Two novel memory circuits were created during this research. The first is a differential latch

cell that uses the MTJs themselves to provide radiation tolerance. The second is a magnetic shadow latch that takes advantage of the bottom gate available in a double-gate silicon-on-insulator technology. The shadow latch concept is extensively studied and optimized. Finally, a new “one-wire” programming method is described. This is a critical aspect of embedded magnetic memory, and limits its energy efficiency. Associated reliability issues are investigated, and various techniques are applied to optimize the circuits.

Acknowledgments

I wish to express my appreciation to my major professor, Dr. Gregory Donohoe, for the considerable efforts he made to accommodate my work circumstances and make it possible for me to complete my degree, and for his continued encouragement throughout it all.

I am indebted to Dr. Yang-Ki Hong at the University of Alabama and Dr. Byoung-Chul Choi at the University of Victoria for their contributions in materials and micromagnetics, which made this research possible.

Dr. David Cox of Apex Microsystems kindly provided material for the analog circuits portion of my qualifying examination. I am grateful to Dr. Cox for many stimulating and educational discussions over the last decade.

This research was supported by the U. S. Air Force Research Laboratory (AFRL) under Grant No. F29601-04-1-206.

Contents

Title **i**

Authorization to Submit Dissertation **ii**

Abstract **iii**

Acknowledgments **v**

Table of Contents **vi**

List of Figures **x**

List of Tables **xiv**

1 Background **1**

1.1 Applications of Embedded Nonvolatile Memory 2

1.1.1 Circumvention 2

1.1.2 Personalization 2

1.1.3 Cryptographic Systems 3

1.2 Ferromagnetism 3

1.3 Electromagnetism 6

1.4 Spin-dependent Tunneling 16

1.5 Magnetic Memories 18

1.5.1 MTJs as memory elements 19
1.5.2 Bulk Magnetic RAM 21
1.5.3 Shape Optimization for MRAM 26
1.6 Radiation Effects in Electronics 29
1.6.1 Total Ionizing Dose Effects 29
1.6.2 Single Event Effects 33
1.6.3 SEU-Resistant Latch Circuits 46

2 Design of a Radiation-Tolerant MTJ Memory Cell 52

2.1 Design Goals 52
2.1.1 Functionality 53
2.1.2 Power Consumption 53
2.1.3 Timing 54
2.1.4 Reliability 54
2.1.5 Radiation Tolerance 55
2.2 Design Constraints 55
2.2.1 MTJ Behavior 55
2.2.2 Reliability 61
2.3 Circuit Design 68
2.3.1 MTJ Simulation Models 68
2.3.2 Prior MTJ Latch Designs 71

3 Differential MTJ Latch 76

4 MTJ Shadow Latch **82**

4.1 Double-gate Transistors 82

4.2 Shadow Latch Circuit 84

4.3 Design Issues 90

4.4 Design Optimization 93

4.4.1 Read Operation With Grounded WRTLN 93

4.4.2 Read Operation With Active WRTLN 113

4.5 Bulk CMOS Implementation 115

5 One-wire Programming **118**

5.1 Overview 118

5.2 Single-Bit Write Line 122

5.3 Serial Write Lines 126

6 Reconfigurable Hardware using MTJs **135**

6.1 MTJs As Programmable Switches 135

6.2 MTJs Applied To Programmable Logic 136

7 Results **139**

Bibliography **140**

Appendices **158**

A Magnetic Field Software Model **158**

B MTJ SPICE Model **165**

C Pairs Plot Script for Gnuplot **168**

D First Shadow Latch Netlist **172**

E Second Shadow Latch Netlist **173**

F Third Shadow Latch Netlist **174**

List of Figures

1.1 Density-of-states for normal and ferromagnetic metals 5

1.2 Field from thin sheet 7

1.3 Field from 3 wires 7

1.4 Field from 9 wires 9

1.5 Horizontal component of magnetic field for 1 μm wide wire of varying thickness 12

1.6 Horizontal component of magnetic field for 2 μm thick wire of varying width . 13

1.7 Horizontal component of magnetic field for wires of varying aspect ratio 14

1.8 Fields produced by clad and unclad lines 15

1.9 Skin depth vs. frequency 16

1.10 Typical MTJ construction 19

1.11 Bulk MRAM crosspoint write 22

1.12 MRAM magnetic field vectors 23

1.13 Typical astroid curve 24

1.14 Bulk MRAM read 25

1.15 Potential free layer shapes 27

1.16 Pac-man shapes 28

1.17 Top view of example transistor 30

1.18 Charge accumulation due to ionizing radiation 32

1.19 Example cosmic particle spectrum 35

1.20 Relative abundance of the elements 36

1.21 Cosmic particle striking a transistor 37

1.22 Generic SOI transistor 38

1.23 CMOS inverter 40

1.24 Cosmic particle strike to inverter pair 42

1.25 Cosmic particle strike to latch 43

1.26 Parasitic bipolar transistors in CMOS 44

1.27 Latch with resistor hardening 47

1.28 DICE schematic 48

1.29 SERT cell schematic 50

2.1 SPICE model of MRR bias dependence 69

2.2 Daughton's latches 71

2.3 Daughton latch (b) simulation with particle strike 72

2.4 Wang latch and Unbalanced Flip Flop 73

2.5 UFF simulation with particle strike 74

3.1 Dual-MTJ latch cell schematic 76

3.2 Simulated cosmic particle strike to PMOS of dual-MTJ cell 79

3.3 MTJ bias during cosmic particle strike 80

4.1 FlexFET cross section 83

4.2 Shadow RAM cell 84

4.3 SPICE simulations 85

4.4 Shadow RAM cell layout 87

4.5 DICE input circuits 89

4.6 Bottom gate small signal gain 92

4.7 Shadow RAM latch simulations replacing MTJs with resistors 94

4.8 Shadow RAM cell layout, second version 96

4.9 Shadow RAM latch simulations replacing MTJs with resistors, second layout . 97

4.10 Simulation detail, second layout 98

4.11 MTJ optimization, second layout 100

4.12 Effect of added capacitance, SPICE data 101

4.13 Effect of added capacitance, waveforms 103

4.14 Shadow RAM cell layout, third version 104

4.15 Pairs Plot, capacitance mismatch and delay skew 106

4.16 Output delay skew vs. capacitance mismatch to WRTLN and capacitance mismatch to drains 108

4.17 Pairs plot, capacitance mismatch and delay skew 109

4.18 Pairs plot, capacitance from Y0-Y3 to VSS and delay skew 110

4.19 Pairs plot, paired capacitance from Y0-Y3 to VSS and delay skew 111

4.20 Shadow latch characterization, third layout 112

4.21 Shadow latch TMRR characterization, third layout 113

4.22 Shadow latch simulation, program/read scenario 114

4.23 Bulk CMOS shadow RAM cell 116

4.24 Bulk CMOS back gate small signal gain 117

5.1 One-wire write scheme 119

5.2 One-wire programming schematic 122

5.3 One-wire programming layout 123

5.4 Write-line current density 124

5.5 Write-line voltage drop 125

5.6 Series connected write line drivers 127

5.7 Optimized series connected write line drivers 129

5.8 N-bit serial one-wire programming schematic 131

5.9 Write energy per bit for series write lines 133

6.1 Programmable logic using MTJ chains 137

A.1 Layout parameters for software model 159

A.2 Geometry of field calculations 161

B.1 MTJ SPICE model schematic 165

C.1 Pairs plot graph example 171

List of Tables

2.1 Material properties for common conductors 68

4.1 Sources of transistor imbalance 91

4.2 Parasitic capacitances in first shadow latch 95

5.1 Typical write line operation 122

5.2 Write line driver corner cases 125

5.3 Serial WRTLN driver, equal V_{DS} 128

5.4 Serial WRTLN driver optimization 129

5.5 Serial WRTLN node voltages, three latches 130

5.6 Serial WRTLN optimization, four latches 130

5.7 Serial WRTLN optimization, five latches 130

6.1 Characteristics of programmable switches 135

C.1 Pairs plot data example 171

Chapter 1

Background

Nonvolatile semiconductor memory has become one of the most significant markets for integrated circuits, driven by consumer electronics such as digital cameras and portable media players. For these applications the motivation for technology development stems from the need for ever greater density. As the number of megapixels in digital cameras climbs, and consumers expect to play video instead of just audio on their hand-held devices, nonvolatile memory needs have quickly moved from megabytes to gigabytes.

In aerospace applications, nonvolatile memory has always played a much different, but vitally important role. A common requirement for these systems is radiation tolerance, so that their reliability is not compromised by the natural radiation environment of space. The daunting physical conditions encountered in aircraft and spacecraft greatly diminish the lifetime of memory devices with moving parts, such as disk drives or tape recorders, and encourage the use of "solid state" semiconductor memory.

This chapter explores the foundations of this research effort, and examines the underlying phenomena. Topics normally encountered in materials science and physics are discussed in order to gain an intuitive understanding of the magnetic tunnel junction and its behavior. We then describe how these devices can be used in practice to store and retrieve information. Finally, an overview of radiation effects in electronics is presented along with some common techniques for their mitigation.

1.1 Applications of Embedded Nonvolatile Memory

There are several unique applications for embedded nonvolatile memory in aerospace systems. They derive from the special functional requirements of these applications as well as the unusually harsh environments in which they must operate.

1.1.1 Circumvention

In situations where electronic systems may be exposed to very short but intense bursts of ionizing radiation, such as nuclear blasts in a space environment, a common problem is that extremely high photocurrents are generated within the integrated circuit. If the voltage drop along the circuit's power supply wiring approaches the nominal supply voltage then the effective supply voltage falls to zero, which is referred to as "rail span collapse" [1]. At higher dose rates the photocurrents may cause permanent damage to power supplies or distribution networks. For some integrated circuit technologies destructive latchup conditions can also occur during a transient dose. A common technique for avoiding transient dose failure is known as circumvention, and involves removing power from the system within a very short period of time, perhaps microseconds, after a nuclear event is detected.

Whether power is lost because of rail-span collapse or circumvention, conventional static RAM cells and flip flops will lose their stored information. Critical data must be stored in some kind of non-volatile memory so that it can be recovered after the transient radiation ends. Desirable characteristics for the non-volatile storage in this application would include a short write time and low sensitivity to total dose effects.

1.1.2 Personalization

The trend in many aerospace applications is toward programmable integrated circuits, such as a Field Programmable Gate Array (FPGA). The cost of designing, manufacturing and

qualifying an integrated circuit can then be amortized over many users. Programmable hardware also allows designers to field systems much more quickly and to adapt to changing requirements. In many cases it is desirable to store the unique programming, or personalization, of a circuit in non-volatile memory so that it begins performing its desired function immediately when power is applied. Non-volatile personalization also saves power in situations where the system is only activated for a brief period of time and then sleeps in an unpowered state until it is needed again.

For personalization applications, the write time of the non-volatile memory is not a critical parameter since writing is infrequent. A more important characteristic is the memory's long term reliability and immunity to radiation effects. Since the memory bits are often scattered amid functional logic blocks it is not practical to use the error-correcting codes that are common in bulk memory.

1.1.3 Cryptographic Systems

Cryptographic systems represent a small but important niche. They typically require a small amount of critical data, such as a key or password, to be stored reliably for long periods of time. If the security of the system has been compromised or if tampering is detected it could be necessary to erase the cryptographic information as quickly as possible and with as little energy as possible.

1.2 Ferromagnetism

One of the foundations of quantum theory is the notion that the electrons in any given atom can exist only at discrete energy levels. For an electron to move to a higher energy level it must absorb a fixed amount, or *quantum*, of energy, and when an electron falls to a lower energy level it releases energy in the same quanta. Every electron in an atom can be

characterized by a unique combination of four quantum numbers which completely describe the energy of the electron. The first three quantum numbers describe the quantization of the electron's total energy, its angular momentum, and the angle between the angular momentum vector and an applied magnetic field [2]. These numbers are integers that, in theory, may have arbitrarily large values. The fourth quantum number describes the electrons "spin", with only two possible values. The electron spin states may be referred to as +½ and –½ or as "up" and "down".

Ferromagnetism is a quantum mechanical phenomenon due to an imbalance in electron spin. In non-magnetic metals the density of states at the Fermi level for spin-up and spin-down electrons is equal, so there is no net magnetization. In ferromagnetic materials the density of states is not equal, and the electrons at the Fermi level will be predominantly spin-up or spin-down [3, 4].

The spin polarization, P, of the electrons in a ferromagnetic material is defined as

$$P = \frac{N\uparrow(E_F) - N\downarrow(E_F)}{N\uparrow(E_F) + N\downarrow(E_F)} \tag{1.1}$$

where $N\uparrow(E_F)$ and $N\downarrow(E_F)$ are the density of state values at the Fermi level for spin-up and spin-down electrons, respectively. An ideal ferromagnetic material would have a spin polarization of 100%, meaning that all of the electrons at the Fermi level had the same spin, but real materials do not exhibit this high degree of polarization. Reported values range from 35% for $Ni_{80}Fe_{20}$ to 90% for CrO_2 [5]. Figure 1.1 illustrates the density of spin states for a normal metal and for a ferromagnetic metal. Since normal metals have an equal density of spin-up and spin-down states at the Fermi level they have a spin polarization of zero.

An applied magnetic field will affect the spin of the electrons in many materials, but in most cases thermal energy will cause the electrons to return to an equal distribution of spin-up and spin-down states after the field is removed. A defining characteristic of ferromagnetic materials is that the electrons in adjacent atoms are quantum mechanically coupled such that

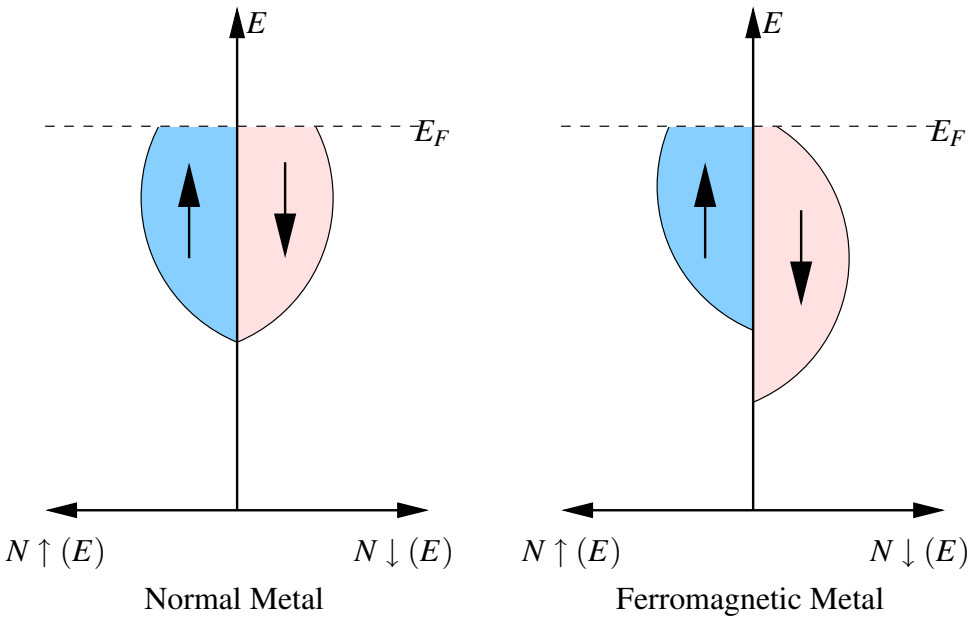

Figure 1.1: Density-of-states for normal and ferromagnetic metals

their spins tend to be aligned, so that a magnetic moment introduced by an external field will cause the ferromagnetic material to retain its magnetization after the field is removed. Another, equally important, class of materials are called antiferromagnetic because the coupling between atoms tends to preserve opposing electron spin between adjacent atoms.

Since ferromagnetism depends on the coupling of electron spin states in a material, there is obviously a fundamental scaling limit for magnetic materials: an isolated atom cannot establish the quantum coupling and therefore cannot exhibit remanent magnetism. In practice, thermal energy can cause spontaneous flips in the magnetic state of particles that are about $10-20$ nm [6]. This is the known as the "superparamagnetic limit" and provides a lower bound on the scaling of magnetic elements. To reliably retain its magnetization for many years, a minimum particle size of approximately 20 nm is required.

1.3 Electromagnetism

To gain an intuitive understanding of the magnetic fields produced by an electric current flowing in a non-ferromagnetic conductor, one can begin with the Biot-Savart law ([7], pp. 232–237). This law specifies the field $d\mathbf{H}$ produced by a current I flowing through a differential length $d\mathbf{L}$ of an infinitesimally thin wire as

$$d\mathbf{H} = \frac{I\, d\mathbf{L} \times \mathbf{a_R}}{4\pi R^2} \tag{1.2}$$

where R is the radius from the wire to the point of interest. If the wire is infinitely long then integrating over $\mathbf{L}$ yields

$$\mathbf{H} = \frac{I}{2\pi r}\mathbf{a}_\Phi \tag{1.3}$$

where $\mathbf{H}$ is a vector and $\mathbf{a}_\Phi$ is the unit vector that obeys the familiar "right-hand rule": if you wrap the fingers of your right hand around a wire and the current flows in the direction pointed by your thumb, then the magnetic field vectors follow the curvature of your fingers. At any point in space the direction of the $\mathbf{H}$ vector is orthogonal to both the direction of current flow and the shortest vector from that point to the wire.

Now consider the field generated by a uniform current I flowing through an ideal conducting sheet, as shown in Fig. 1.2. The sheet has a finite width of $3Y$ but infinite length into the paper. We are interested in determining the horizontal (x) component of the magnetic field (H_{ZX}) at the point labeled Z at a distance Y above the center of the sheet.

Suppose the sheet is modeled as three wires, labeled A, B and C in Fig. 1.3. These ideal wires are symbolized as crosses, indicating that the current they carry flows into the plane of the paper. The wires are spaced apart by a distance Y, and each wire now carries one third of the total current, I.

The magnitude of the magnetic field generated at any point by a current in an infinite ideal

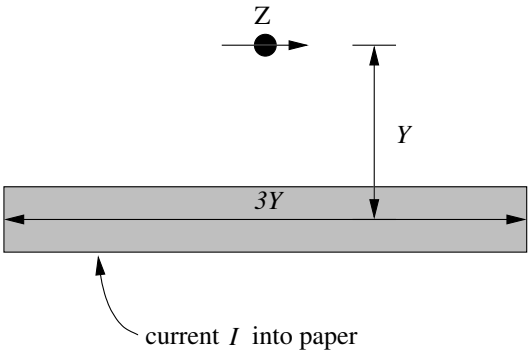

Figure 1.2: Field from thin sheet

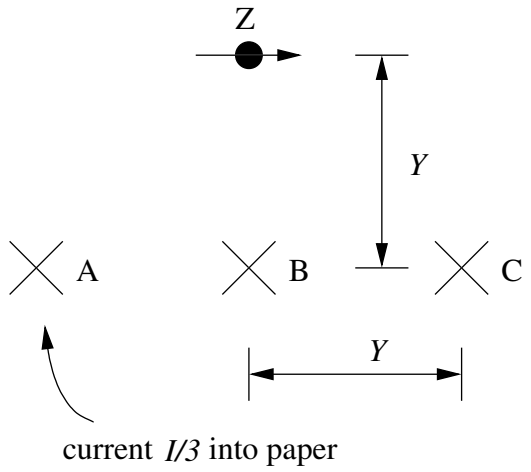

Figure 1.3: Field from 3 wires

wire is proportional to the current and to the distance from the wire. We define the total field produced at Z by wire B as H_B, and since wire B is directly below Z this field vector points directly to the right. Therefore, the x-component of this field vector, H_{BX}, is exactly equal to H_B. For the remainder of the analysis all fields will be normalized to the field vector H_{BX} that results from a current of $I/3$, which is referred to as H_0.

Now the distance from wire A to Z is $Y\sqrt{2}$, so the field generated at Z by wire A is

$$H_A = \frac{H_B}{\sqrt{2}} \tag{1.4}$$

and the x component of the field vector is

$$H_{AX} = \frac{1}{\sqrt{2}} \times \frac{H_B}{\sqrt{2}} \tag{1.5}$$
$$= \frac{H_B}{2} \tag{1.6}$$
$$= \frac{H_{BX}}{2} \tag{1.7}$$
$$= \frac{H_0}{2} \tag{1.8}$$

Because of symmetry $H_{AX} = H_{CX}$ so with superposition

$$H_{ZX} = H_0 + \frac{H_0}{2} + \frac{H_0}{2} \tag{1.9}$$
$$= 2\,H_0 \tag{1.10}$$

If we replace the thin conducting sheet with a sheet that is 3 times thicker, this can be represented by the wires shown in Fig. 1.4. Now there are nine wires, each carrying a current equal to $I/9$. Since the geometry of wire B with respect to point Z is unchanged it is clear that now

$$H_{BX} = \frac{H_0}{3} \tag{1.11}$$

Now, wire E is also directly below point Z but at a distance of $2Y$, so

$$H_{EX} = \frac{H_{BX}}{2} \tag{1.12}$$
$$= \frac{H_0}{6} \tag{1.13}$$

Similarly, since wire H is $3Y$ from point Z,

$$H_{HX} = \frac{H_{BX}}{3} \tag{1.14}$$
$$= \frac{H_0}{9} \tag{1.15}$$

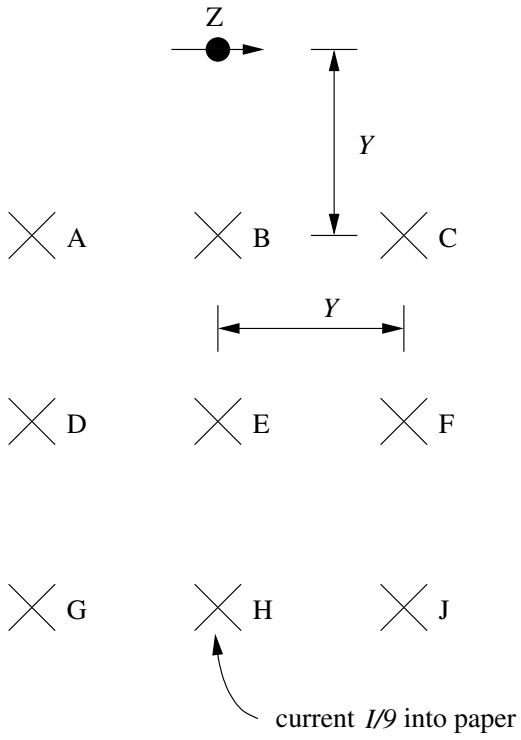

Figure 1.4: Field from 9 wires

So by splitting the original wire B in Fig. 1.3 into three wires (B, E, and H) in Fig. 1.4 and reducing the current density by a factor of three, the contribution to H_{ZX} decreased only to $\frac{11}{18}H_0$, or about 61% of the original value. In this case, where the field vector of interest is orthogonal to a line that passes through the point of interest and all of the wires, the contribution of the wire that is nY away from point Z will be H_{BX}/n. If there are m such wires, each carrying a current of I/m, then the field generated by them at Z is

$$\frac{H_0}{m}\sum_{n=1}^{m}\frac{1}{n} \tag{1.16}$$

where the summation is the familiar harmonic sequence. For example, if $Y = 25$ nm and a single wire is replaced with 24 wires (extending 600 nm below Z), the current density is reduced to 4.2% of the original value but H_{ZX} is still 15.7% of its original value.

Turning to the other wires, wire A is still $Y\sqrt{2}$ from point Z but it is now carrying a current of

$I/9$ so the generated field is reduced by a factor of 3 and $H_{AX} = H_0/6$. Wire D is $Y\sqrt{5}$ from Z and the x component of its field is $2/\sqrt{5}$ of the total, so

$$H_{DX} = \frac{1}{\sqrt{5}} \times \frac{2}{\sqrt{5}} \times H_{BX} \tag{1.17}$$

$$= \frac{2}{5} \times \frac{H_0}{3} \tag{1.18}$$

$$= \frac{2\,H_0}{15} \tag{1.19}$$

Wire G is $Y\sqrt{10}$ from Z, with fractional x component of $3/\sqrt{10}$, so

$$H_{GX} = \frac{1}{\sqrt{10}} \times \frac{3}{\sqrt{10}} \times H_{BX} \tag{1.20}$$

$$= \frac{3}{10} \times \frac{H_0}{3} \tag{1.21}$$

$$= \frac{H_0}{10} \tag{1.22}$$

Adding the contributions of wires A, D, and G results in a total of $\frac{12}{30}H_0$, which is 80% of the value of H_{AX} in Fig. 1.3 even though the current density has been reduced by a factor of 3. For wires that are farther away from Z in the x dimension there will be an even smaller reduction in the field resulting from making the sheet thicker, since the total distance from the wire to point Z changes by a smaller fraction.

Again using symmetry and superposition, the total value of H_{ZX} is

$$H_{ZX} = \frac{11}{18}H_0 + \frac{12}{30}H_0 + \frac{12}{30}H_0 \tag{1.23}$$

$$= 1.411\ H_0 \tag{1.24}$$

Therefore, by replacing the three wires of Fig. 1.3 with the nine wires of Fig. 1.4 the current density was reduced by a factor of 3 but H_{ZX} remains at about 71% of its original value. Conversely, if the total current is increased to achieve the original value of H_{ZX} the current

density would still be less than half of its original value.

Repeating this exercise with a wider sheet, particularly one much wider than the spacing Y, would show little reduction in the field as the sheet is made several times thicker. However, as the thickness of the sheet becomes a significant fraction of its width it no longer behaves like an ideal sheet and the field falls off dramatically near the edges.

A software program that models the magnetic field above a rectangular conductor was created to extend the concepts discussed above (see Appendix A). It allows the conductor to be treated as if it were thousands or millions of individual wires and computes the horizontal component of the field at any desired number of points above the conductor. The model assumes that the conductor and surrounding material are isotropic and have a relative permeability of 1.0. Input parameters to the model include the total current in the conductor, the conductor dimensions, the dimensions of the area of interest above the conductor, and the granularity (grid size) of the model. The output is a table of coordinates for points in the area of interest, along with the value of the horizontal component of the magnetic field, in Oersteds.

Fig. 1.5 illustrates how the magnetic field just above the surface of the wire changes as the wire thickness varies. Each curve shows the strength of the horizontal (x) component of the magnetic field approximately 2.5 nm above the surface of a wire that is 1 μm wide. This data was created using a grid size of 1 nm. As expected, the uniformity of the field decreases as the wire thickness becomes larger and the wire no longer behaves like an ideal sheet. However, the current density decreases much more rapidly than the average field strength so there is a clear motivation to increase the wire thickness.

Fig. 1.6 is also obtained from the software model, but in this case the wire thickness is fixed at 2 μm and the wire width varies. Only the portion of the field above the center 1 μm of the wire is shown. In this case the magnitude of the field falls nearly as fast as the current density when the wire is widened, so there is much less benefit in this case.

Another way to consider these factors is to compare the fields produced by two wires with the

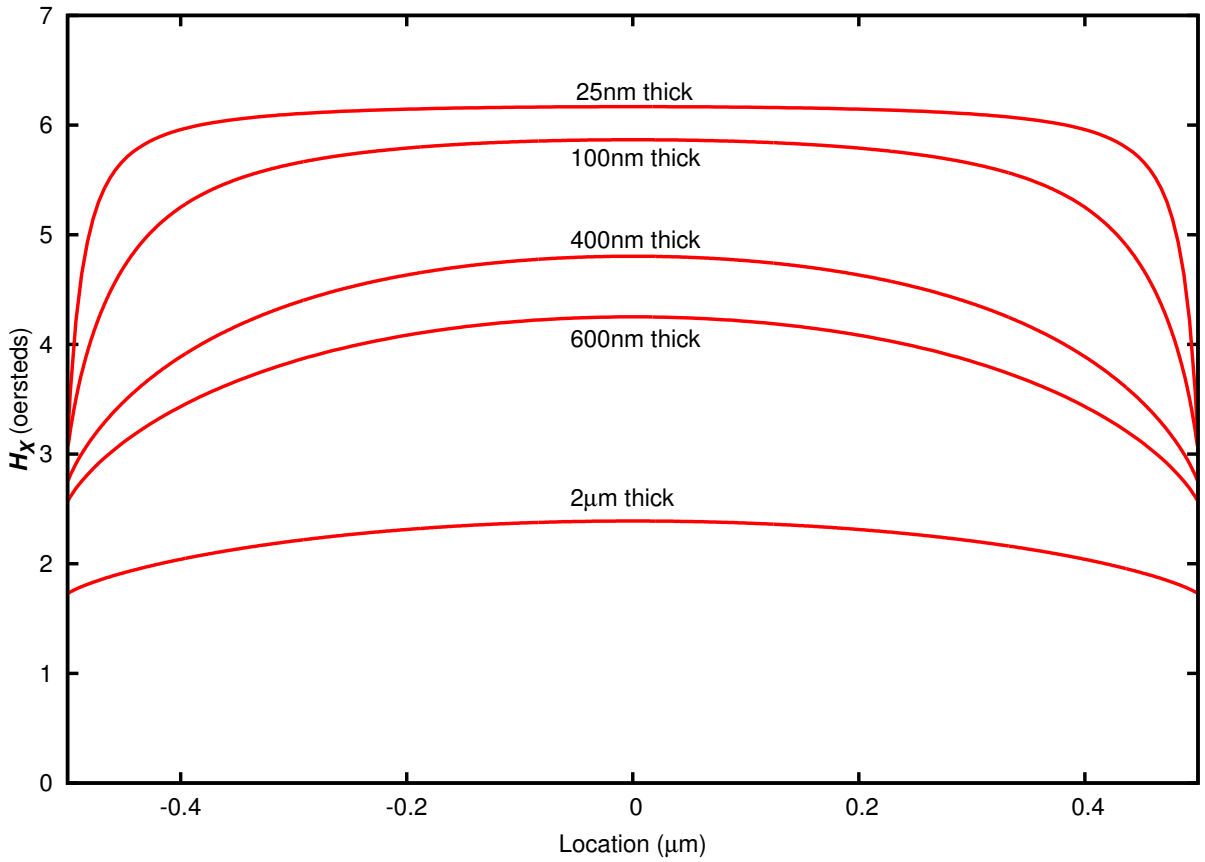

Figure 1.5: Horizontal component of magnetic field for 1 µm wide wire of varying thickness, per mA of current, 2.5 nm above the surface

same cross-section area (two square microns) but different aspect ratios, as shown in Fig. 1.7. The top graph is for a wire that is 2 µm wide and 1 µm thick while the bottom graph is for a wire that is 1 µm wide and 2 µm thick. The origin of coordinates for this graph is placed at the center of the top surface of the conductor, so the Y axis represents the height above the conductor. The current density is the same in both cases. As expected, the narrower wire produces a more intense but less uniform field. At a height of 5 nm above the center of the conductor, the 1 µm wide wire has a peak horizontal field of 2.83 Oe while the 2 µm wide wire has a peak of 2.26 Oe, both carrying 1 mA of current. (Note that the colormap range in this figure is compressed.)

The analyses presented above have assumed that the materials used in the conductor, as well as the surrounding dielectric, have a relative permeability of one. However, it is possible to

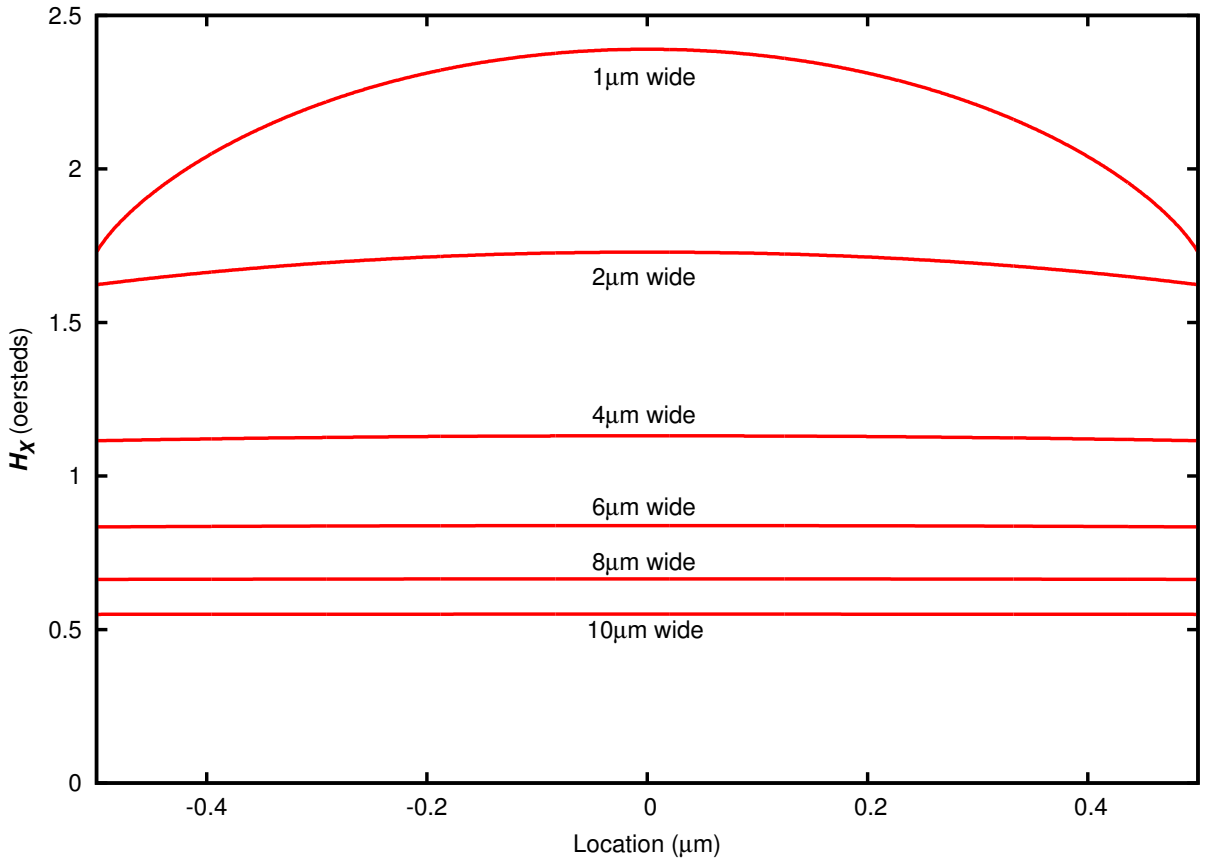

Figure 1.6: Horizontal component of magnetic field for 2 μm thick wire of varying width, per mA of current, 2.5 nm above the surface

introduce material layers with a high permeability in order to improve the efficiency of converting the electrical current into a magnetic field within the free layer. For example, if the MTJ free layer is positioned above the conductor, and the other three sides of the conductor are clad with a 200 nm thick layer of a material with a relative permeability, μ_R, of 1000 then the field shown in Fig. 1.8 will result. (The data for this graph was generated by QuickField, a commercial modeling program [8].) These plots are for metal conductors that are 2.0 μm wide and 1.0 μm thick, carrying a current of 1.0 mA ($J = 5 \times 10^8$ A/m^2). Similar structures have been fabricated and tested with a 50 nm cladding, and the effective field strength increased by a factor of about 1.8 [9].

All of the magnetic field models discussed above have assumed a direct current and ignored high frequency effects, and at this point the validity of that assumption should be examined.

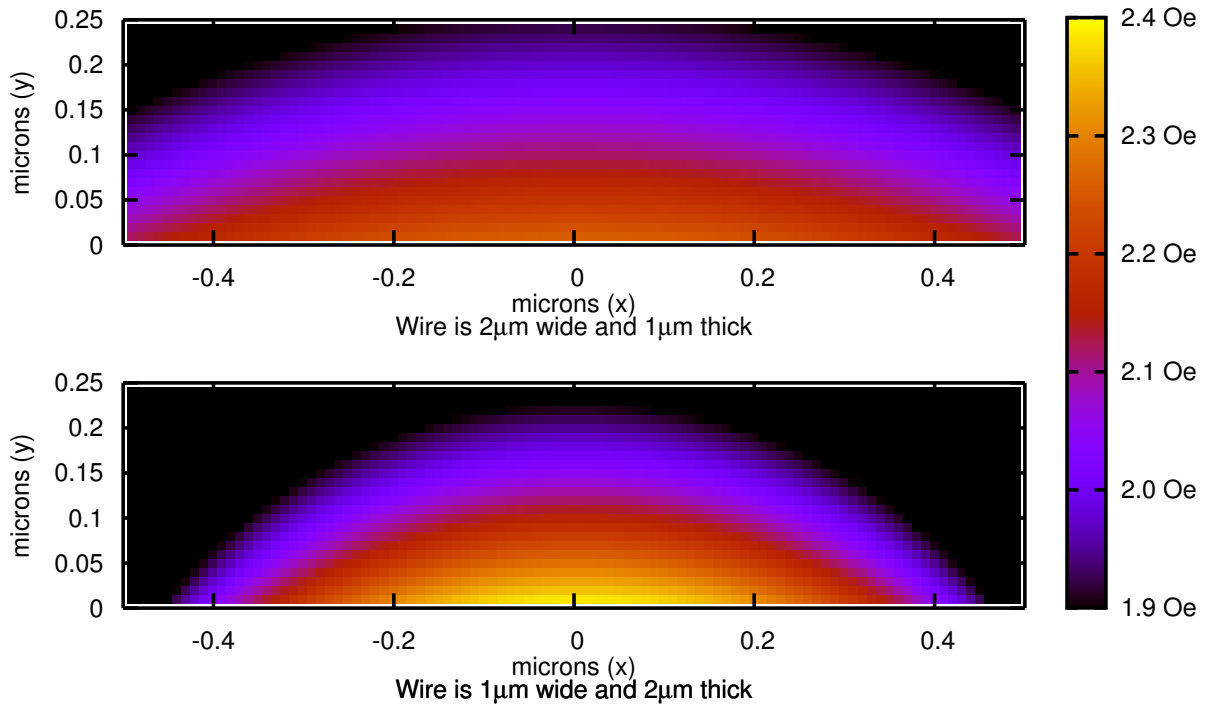

Figure 1.7: Horizontal component of magnetic field for wires of varying aspect ratio, per mA of current

At d.c. the current density is uniform over the cross section of the conductor because the resistivity is also uniform and inductance can be neglected. At high frequencies the self inductance of the conductor becomes significant and non-uniform. If the conductor is visualized as a collection of thin filaments, it is clear that the filaments in the center of the conductor will experience the largest magnetic field coupling to surrounding filaments and therefore have the greatest inductance, while the filaments closer to the surface have less magnetic coupling and lower inductance. Since the impedance at the core of the conductor is higher, the current density will tend to decrease there and increase near the surface, a phenomenon known as the *skin effect*.

A useful metric for the skin effect is the *skin depth*. The a.c. current density decreases exponentially when moving from the surface into the conductor, and the skin depth is defined as that point where the current density falls to e^{-1} times the value at the surface. This

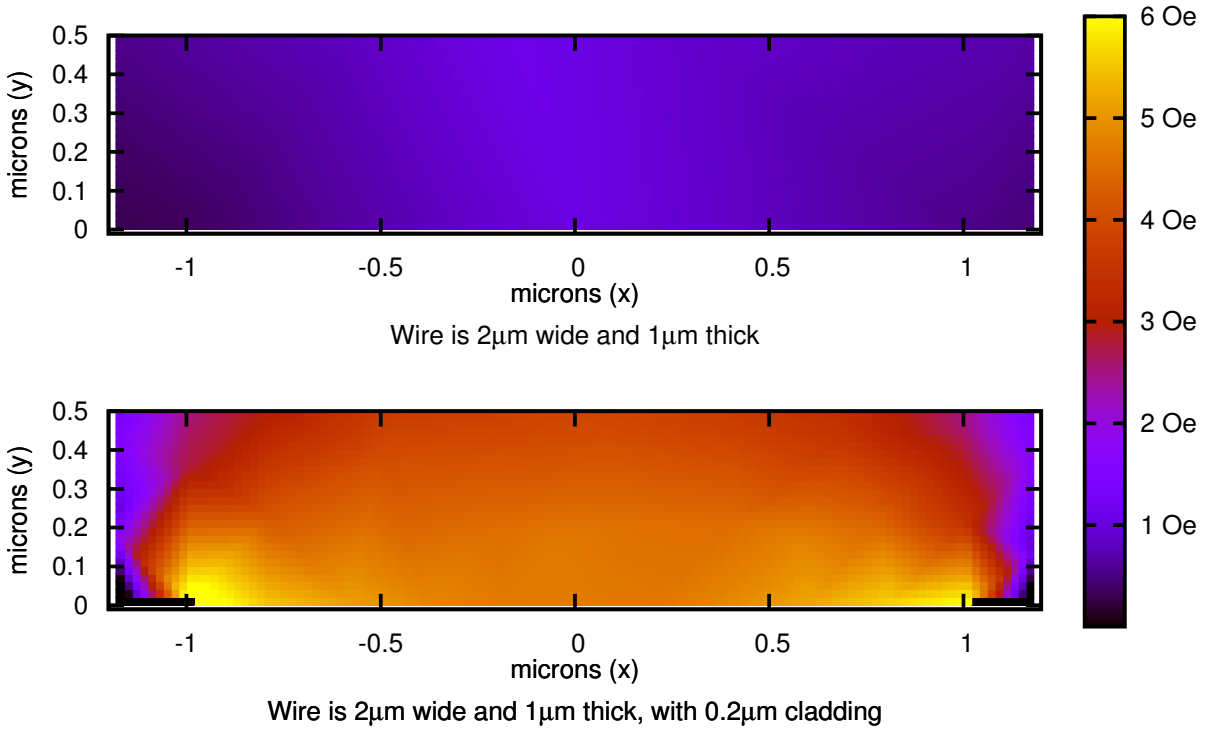

Figure 1.8: Fields produced by clad and unclad lines

parameter can be calculated as

$$\delta = \sqrt{\frac{\rho}{\pi f \mu}} \tag{1.25}$$

where δ is the skin depth, ρ is resistivity of the conductor, f is frequency, and $\mu = \mu_0 \mu_R$ is the permeability of the conductor [7]. Figure 1.9 shows the skin depth for aluminum, copper, and gold conductors from 1 GHz to 10 GHz, which corresponds to pulse rise times from 1 ns to 0.1 ns. For rise times slower than 0.2 ns and conductor thicknesses less than 1 µm, the skin effect does not cause appreciable non-uniformity in the current density. Similarly, inductance should not significantly increase the impedance of the write line for the conductor dimensions and frequencies of interest here [10].

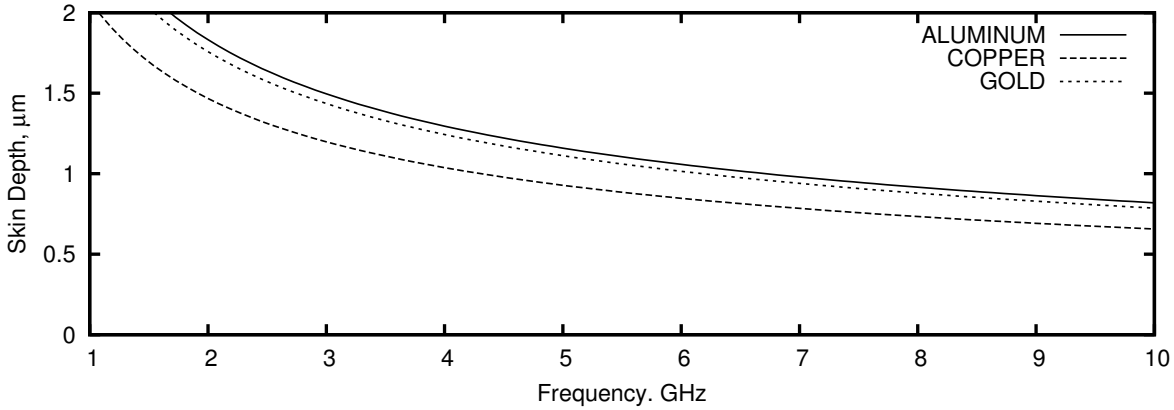

Figure 1.9: Skin depth vs. frequency

1.4 Spin-dependent Tunneling

Electron *tunneling* is another example of quantum mechanical principles, and one that contradicts conventional wisdom based on classical mechanics. Quantum theory states that when an electron is on one side of a high, but very thin, energy barrier there is a significant probability that the electron will appear on the other side of the barrier by tunneling through it without gaining sufficient energy to surmount it [11, 12].

The Pauli exclusion principle applies to tunneling, meaning that an electron will pass through the energy barrier only if a vacancy with the same quantum state is available on the other side. Thus, an electron with a given spin state can only tunnel to a vacancy of the same spin. When the tunnel barrier separates two non-ferrous metals, with no significant spin polarization, there is no shortage of vacancies with either spin state. However, by separating two ferromagnetic materials with an extremely thin insulating layer it is possible to exploit the Pauli exclusion principle to detect the relative alignment of the magnetization of ferromagnetic layers. When the magnetization vectors of the two ferromagnetic layers are parallel then they share the same spin polarization and tunneling will occur more readily than if the magnetization vectors are antiparallel. This assembly of two conducting layers separated by a thin insulator is referred to as a *tunnel junction*.

Spin polarization and spin-dependent tunneling were observed in the early 1970's but only in superconductors at cryogenic temperatures [13–15]. Julliere studied tunnel junctions formed with two ferromagnetic layers separated by an insulator, but demonstrated a significant effect only at very low temperature [16]. Two decades would pass before improvements in materials would allow researchers to demonstrate non-trivial levels of spin-dependent tunneling at room temperature [17–19]. Once the potential usefulness of a magnetic tunnel junction (MTJ) was realized, research efforts in this area expanded dramatically.

The change in tunneling currents between the parallel and antiparallel cases is measured as a macroscopic change in the tunneling magnetoresistance (TMR) of the tunnel junction. This resistance difference, the tunneling magnetoresistance ratio (TMRR), is a critical parameter for determining the usefulness of a tunnel junction. The relationship between spin polarization and TMRR was expressed by Julliere [16] as

$$TMRR = \frac{R_{AP} - R_P}{R_P} = \frac{2P_1P_2}{1 + (P_1P_2)} \tag{1.26}$$

where P_1 and P_2 are the spin polarization values for spin-up and spin-down electrons, R_{AP} is the resistance of the junction when the magnetic fields have antiparallel orientation and R_P is the resistance when the fields are parallel. Note that by defining the TMRR with R_P, the smaller resistance, in the denominator it is possible to achieve TMRR values greater than 100%, and this is typical in recent literature [6]

Tunneling magnetoresistance is not just a function of the spin polarization of the ferromagnetic electrodes, but is also strongly influenced by the insulating layer and the interfaces between layers [20]. Interactions at the interface between a ferromagnetic layer and an insulating layer can flip the spin of an electron, and these electrons can tunnel to minority spin states in the other ferromagnetic layer, resulting in a spin-flip component of tunneling

current. Incorporating this factor into Julliere's model gives

$$-\frac{\Delta R}{R} = \frac{2(1-\gamma)P_1P_2}{1+P_1P_2+\gamma(1-P_1P_2)} \tag{1.27}$$

where $0 \leq \gamma < 1$ characterizes the spin-flip effect at the interface [21]. This equation is equivalent to Julliere's when $\gamma = 0$, meaning that no spin-flip tunneling occurs.

1.5 Magnetic Memories

Magnetism has been used as a means of storing information since the earliest days of electronic computers. The original computer "core" memory consisted of tiny ferromagnetic toroids that were threaded with fine wires in a two-dimensional array. The orientation of the field in a particular toroid was set by simultaneously driving a current along the horizontal (row) and vertical (column) wires that intersected at the location of the toroid. The state of a particular toroid was determined by attempting to set the direction of its field and observing whether or not a current pulse occurred. If a pulse did occur then the toroid had flipped its orientation, if no pulse was observed then the toroid was already oriented in the programmed direction. This destructive read-out technique complicated applications for core memory.

Later, data bits were encoded as the orientation of the magnetic domains in a bulk ferromagnetic medium in the shape of a cylinder, wire, or disk. Information was retrieved (non-destructively) by sensing the current induced when a coil moved across the medium. Although this type of magnetic memory is much easier to manufacture it has the disadvantage of requiring moving parts, and it is difficult to quickly read or write data that is not physically colocated on the medium.

1.5.1 MTJs as memory elements

With the development of MTJs that exhibited high room-temperature magnetoresistance ratios, it became possible to envision memory devices that combine the random access capability of the old core memory with modern integrated circuit fabrication technology [22]. Such a magnetic random access memory, or MRAM, offers several advantages over the commonly used floating-gate non-volatile memory. An MRAM cell can be written much faster, using less energy, than a floating-gate cell. MRAMs appear to be readily scalable to decananometer integrated circuit technology without sacrificing reliability. They also should be less susceptible to the effects of ionizing radiation, an important factor for aerospace and military applications.

In the years after practical MTJs were demonstrated in the laboratory, considerable effort was devoted to the manufacturing and circuit design issues that would lead to MRAM products [23–25]. By the year 2000 several working (albeit quite small) MRAMs were reported [26, 27].

The MTJs used in a typical MRAM share a general form of construction and operation, as shown in Fig. 1.10. Conceptually, the MTJ consists of two ferromagnetic layers separated by

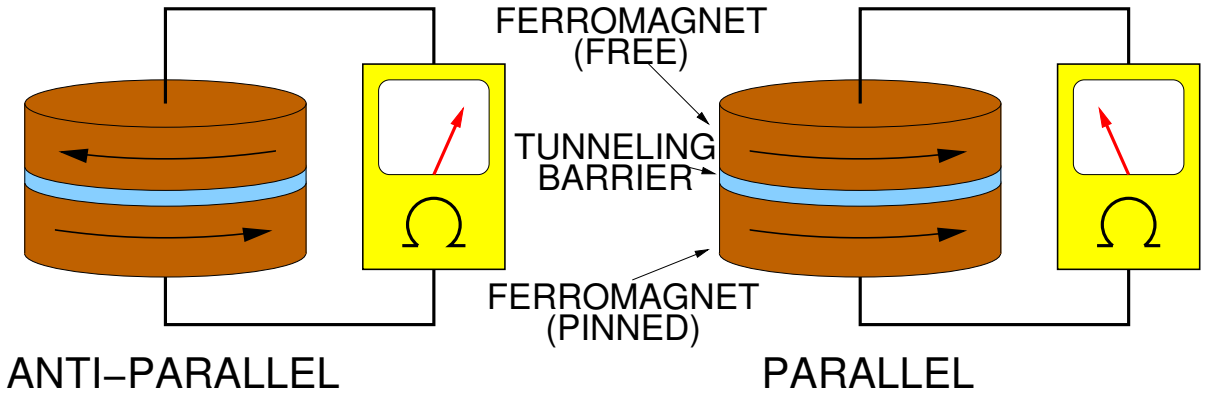

Figure 1.10: Typical MTJ construction

a thin tunneling barrier. The bottom ferromagnetic layer is “pinned”, meaning that the orientation of its magnetic field is fixed during manufacturing and has a relatively high

coercivity. The top ferromagnetic layer is “free” and has a relatively low coercivity so its magnetic field orientation can be readily reversed by a small applied field. The resistance of the MTJ is measured to determine whether the retained fields are parallel or antiparallel, which in turn indicates the logic state of the data bit stored in the cell.

The pinned layer is typically a layered structure itself, comprising various combinations of ferromagnetic, metallic, and antiferromagnetic materials [28]. The layer in contact with the tunneling barrier is a ferromagnetic layer, and in the simplest case this ferromagnetic layer is then in contact with an antiferromagnetic layer [29]. Exchange-biasing between these two layers tends to lock the orientation of the ferromagnetic layer. It is also possible to create a synthetic antiferromagnet by using two ferromagnetic layers separated by a ruthenium spacer layer, where the ferromagnetic layers have opposite magnetic orientations that are preserved by flux coupling between them [30]. More elaborate stacks combine natural and synthetic antiferromagnetic layers [31, 32].

The free layer is commonly a single ferromagnetic layer. Early work was done with permalloy ($Ni_{80}Fe_{20}$) because thin films would retain the magnetic behavior of the bulk material and it could be given an easy axis by applying a magnetic field during heat treatment [30]. Although permalloy continues to be used, a broad assortment of other materials have been studied, including Co [33], CoFe [34], and NiFeCo [29, 35, 36]. As the size of the free layer is reduced the required switching field increases, which would limit the scalability of MTJs [34, 36]. However, using a synthetic antiferromagnet for the free layer instead of a single layer can reduce this effect [37]. The thickness of the free layer is also a factor, with the general observation being that for layers thinner than 5 nm, thicker layers require a stronger switching field [38]. Later studies of elliptical shapes found that this was true for thicknesses up to some critical value, which was in the range of 24 nm to 40 nm, but the required field decreased for thicker layers [39]. The patterned shape of the free layer also has a strong influence on its properties, as is discussed in Section 1.5.3.

Since the tunneling insulator is grown on top of the bottom electrode (either the free layer or the pinned layer) in the MTJ, the microstructural characteristics of this electrode are critical to the quality and reliability of the resulting junction. In a study of MTJs where the roughness of the electrode surface was controlled, Kim et al. reported that a rough surface would reduce the resistance-area (RA) product and increase the bias dependence of TMRR [40, 41]. As the surface roughness approached the thickness of the barrier, the RA product fell by 60% and the breakdown voltage of the barrier also decreased.

The tunneling barrier is perhaps the greatest manufacturing challenge for high-density MRAM, as it is this layer that determines most of the important characteristics of the MTJ. The first MTJs typically used Al_2O_3 as the insulating layer because it can be grown readily by depositing a very thin aluminum film (typically $2-3$ nm) and then oxidizing the aluminum. It is difficult to achieve pinhole-free layers thinner than about 1 nm, and layers thicker than about 3 nm tend to be incompletely oxidized [42]. While Al_2O_3 is an amorphous insulator, MgO is crystalline and causes less scattering of the tunnel electrons, resulting in significantly larger TMR ratios (see also the discussion in Section 2.2.1).

1.5.2 Bulk Magnetic RAM

A typical MRAM shares the row/column addressing scheme used in early core memories, but the conductors are now patterned as thin metal films on an integrated circuit rather than fine wires threaded manually through the center of a toroid.

Write Operation

Changing the magnetic orientation of the free layer in a single MTJ is accomplished by creating two orthogonal fields whose vector sum is sufficient to reverse the magnetization. Suppose that the “bit line” in the Fig. 1.11 is one of many parallel conductors that run in one direction across the chip, and the “word line” is one of many parallel conductors that also run

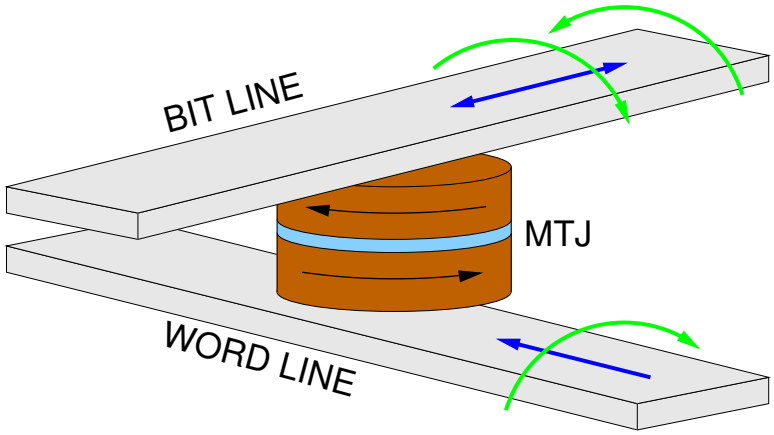

Figure 1.11: Bulk MRAM crosspoint write

across the chip but in an orthogonal direction to the bit lines. (For convenience we arbitrarily refer to the bit lines as vertical columns and to the word lines as horizontal rows.) Suppose that a current is forced through only one bit line and only one word line, while all other bit and word lines are grounded. All of the MTJs under that bit line will experience the field produced by the bit line current, and all of the MTJs above that word line will experience the field produced by the word line current. The single MTJ at the intersection, or crosspoint, of the selected bit and word lines experiences the sum of both the bit line and word line fields.

The example of Fig. 1.11 assumes that the current (the blue arrow) in the word line always flows in the same direction, producing a magnetic field that will always have the same orientation (the green arc). The familiar "right-hand rule" reminds us that if the thumb of the right hand is pointed in the direction of current flow then the resulting magnetic field is oriented like the fingers curled toward the palm. As suggested in the figure, the current direction in the bit line can be reversed to change the orientation of the magnetic field applied to the MTJ. Fortunately, applying a hard axis field helps to change the magnetization reversal from domain wall movement to coherent reversal, thereby reducing the required field along the easy axis [36, 38].

It is important to note that Fig. 1.11 is not drawn to scale, and in particular that the MTJ free layer is much thinner than suggested by the drawing. For practical purposes we can ignore

magnetic fields that are not coplanar with the ferromagnetic film, and apply the results from Sec. 1.3. Fig. 1.12 illustrates how the fields from the bit and word lines combine to produce a net field. In an MRAM the two generated fields are always orthogonal, so the magnitude and

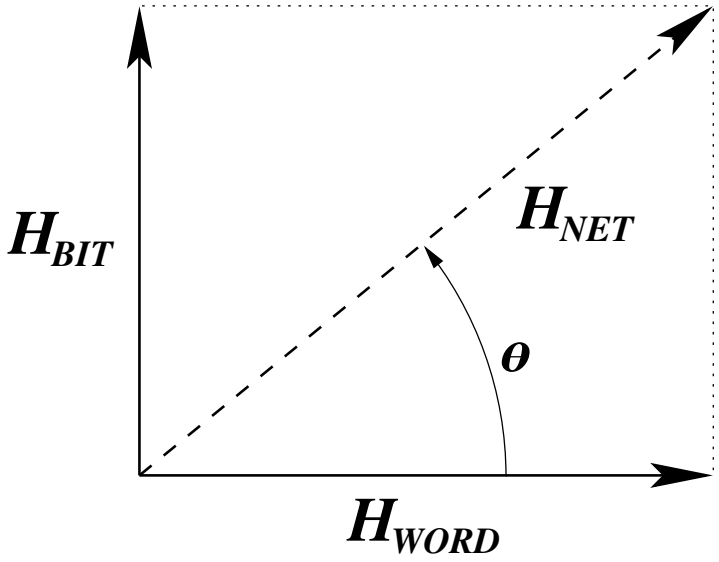

Figure 1.12: MRAM magnetic field vectors

angle of the net resulting field is

$$H_{NET} = \sqrt{H_{BIT}^2 + H_{WORD}^2} \tag{1.28}$$

$$\Theta = \arctan \frac{H_{BIT}}{H_{WORD}} \tag{1.29}$$

For example, if we assume that the currents in the bit and word lines are equal and that they are equidistant from the free layer, then clearly the resulting magnetic field generated in the plane of the free layer will have a magnitude equal to $\sqrt{2}$ times the magnitude of either field alone, and it will be oriented 45° from either of the lines. Note that the vector sum of the fields is almost 30% less than their arithmetic sum, indicating that this portion of the in-plane field (and the energy required to create it) has been canceled and provides no benefit.

However, the ability to control the direction and magnitude of the net magnetic field during

the write operation provides the solution to the *selectivity* problem in MRAMs, and enables the tight packing of MTJs in a row/column arrangement. Recall that all of the unselected MTJs on the selected bit and word lines (i.e. those MTJs not at the intersection of the two lines) still experience the field produced by either the bit line or the word line alone. It is critical that these "half-select" fields do not disturb the data stored in an MTJ even though they are stronger than the field experienced by the single selected MTJ. The MTJ's free layer must therefore have an anisotropic response to magnetic fields, which motivates the research into magnetic elements that have the desired shape anisotropy. This characteristic is often

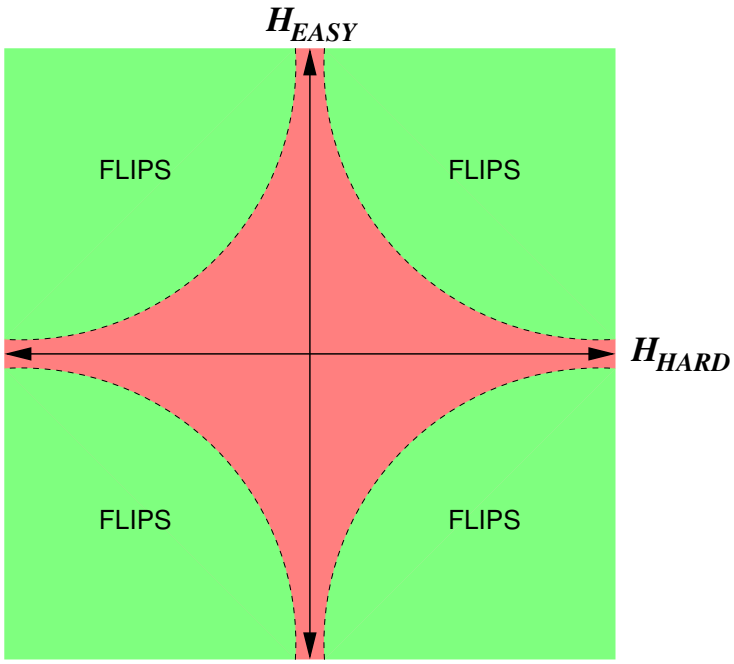

Figure 1.13: Typical astroid curve

described in terms of an "astroid" [1] diagram, like the idealized one shown in Fig. 1.13. Note that very large values of H_{EASY} or H_{HARD} alone will not flip the state of the bit, but the combined effect of moderate fields applied simultaneously in both directions can accomplish the task.

[1] In geometry, an astroid is the figure described by the parametric equations $x = \cos^3\theta$ and $y = \sin^3\theta$.

Read Operation

An example of how data can be read from an MRAM cell is shown in Fig. 1.14. The word line is omitted from this figure; instead the "digit line" is shown. The digit line is parallel to the word line but electrically and functionally distinct. Since it is not involved in writing the MTJ the digit line may be more distant from the MTJ and need not carry a large current. The

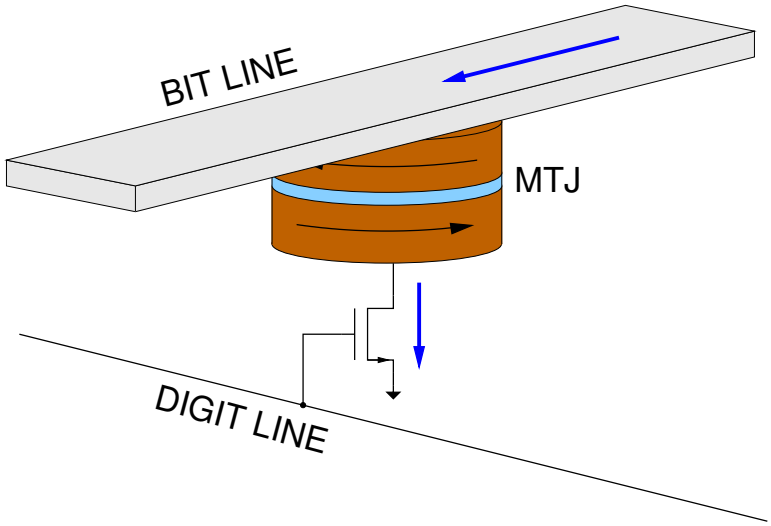

Figure 1.14: Bulk MRAM read

MTJ is assumed to be in ohmic contact with the bit line, and the other side of the tunneling junction is connected to the drain of an NMOS transistor. (There is no electrical connection to the word line.) We desire to measure the resistance through a single MTJ that is located at the intersection between a specific bit line (column) and a specific digit line (row). This can be accomplished by applying a positive voltage to the bit line and digit line in question and grounding all other bit lines and digit lines. The high voltage on the selected digit line enables all of the NMOS transistors associated with the MTJs in that row, but the NMOS transistors in all other rows are disabled and no current can flow through their MTJs. Furthermore, no current will flow through the MTJs attached to grounded bit lines since there is no voltage across their junctions. Thus, the current flowing in the bit line will be the tunneling current through the selected MTJ.

1.5.3 Shape Optimization for MRAM

The shape of the free layer is very important to insure fast, symmetrical, and reliable reorientation of its field. Very small magnetic elements rely on shape anisotropy to result in two stable orientations of the magnetization with very good alignment to the pinned layer in both the parallel and antiparallel states, as well as immunity to spontaneous reorientation due to thermal energy. On the other hand, excessive anisotropy can require stronger fields to switch the free layer, increasing power consumption.

The free layer shape can also affect how its magnetization is reoriented, whether by *coherent reversal* or *domain wall movement*. If the magnetization of the element exists as a single domain, and the orientation of this domain simply rotates to align with an applied field, then the reorientation is said to occur by coherent reversal. This is the preferred mechanism, as it is allows for sub-nanosecond switching and prevents ringing [43]. If coherent reversal does not occur, then the reorientation proceeds through domain wall motion. The applied field forms small new domains at the edges of the element, aligned 180° from the magnetization of the rest of the structure. These domains grow as their walls expand, eventually merging into a single domain with the new orientation. In some cases, domain wall motion will result in the undesirable creation of circular, or "vortex", domains that require much higher fields to annihilate.

Consequently, a great deal of research has focused on finding the most suitable shape for the free layer, including those shown in Fig. 1.15. For MRAM applications, a desirable shape is one that switches reliably, quickly, and with a small applied field along the easy and hard axes. At the same time, the shape should not be affected by relatively large fields along either the hard or easy axis alone, in order to provide selectivity. To be commercially viable the shape must be easily created with the standard lithographic techniques used in integrated circuit manufacturing.

The earliest thin film magnetic elements were simple rectangles, but edge domains can cause

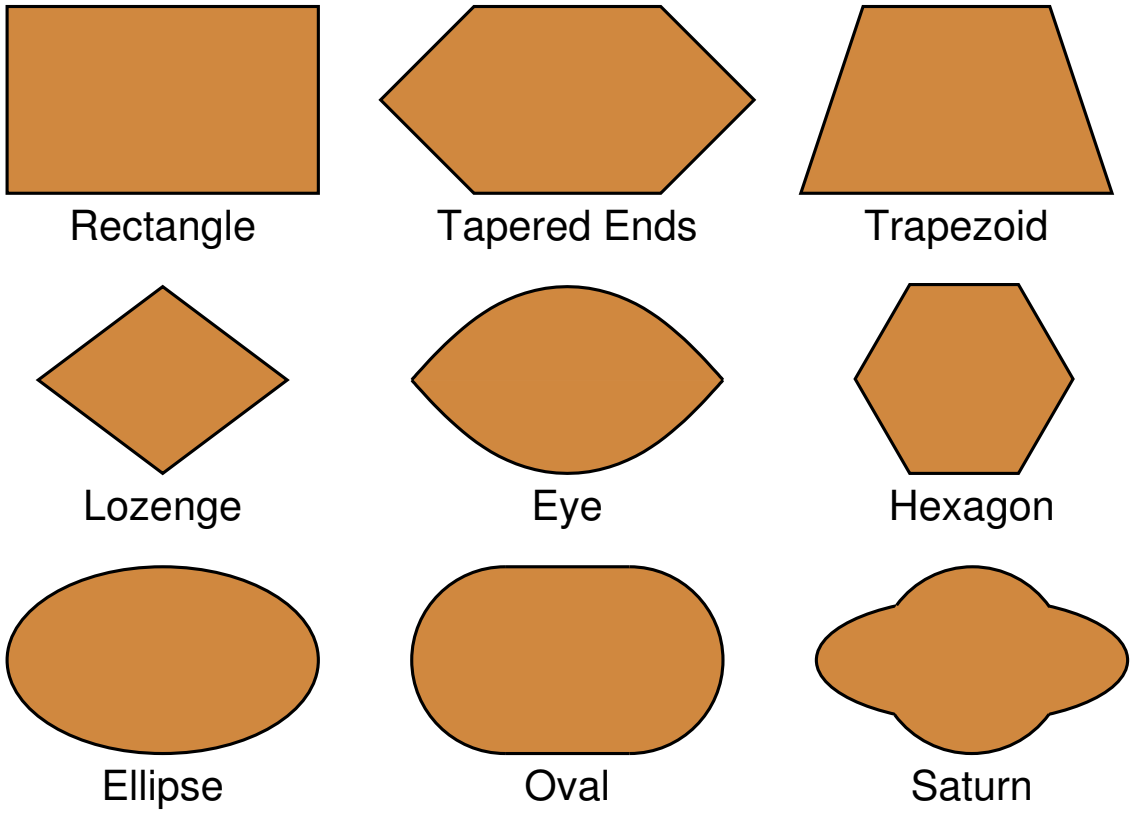

Figure 1.15: Potential free layer shapes

the switching characteristics to change depending on the remanent state, increasing the required switching field by 16% [33]. Rectangles with tapered ends were studied, but they tended to form a variety of domain configurations and were very sensitive to the shape of the tips [44]. Variations on this theme included the hexagonal [29] and "lozenge" shapes [34, 44]. Trapezoids have also been studied, but they appear to be only of academic interest when studying domain wall motion [36].

Non-rectilinear shapes have also been proposed. A circle offers no shape anisotropy but it can be stretched in one dimension to form an oval or ellipse [38, 45, 46]. The "eye" shape is like an ellipse with pointed ends and is formed from two facing parabolas [34, 38]. Combining two of these simple shapes can produce a more exotic design, such as elements that resemble the planet Saturn [47, 48].

Obviously, some of the studied shapes are actually the same underlying geometry with a varying aspect ratio (i.e. length-to-width ratio), and it should come as no surprise that the aspect ratio is another variable with a significant impact on the element's behavior. For ellipses with aspect ratios of about five and higher, the magnetization tends to be uniformly aligned with the long axis of the element because of shape anisotropy, while lower aspect ratios tend to have to more complex remanent arrangements of multiple domains [39].

A novel shape that has been studied as part of this research is the pac-man (PM) element, shown in Fig. 1.16, which is named for the video game character that it resembles. The

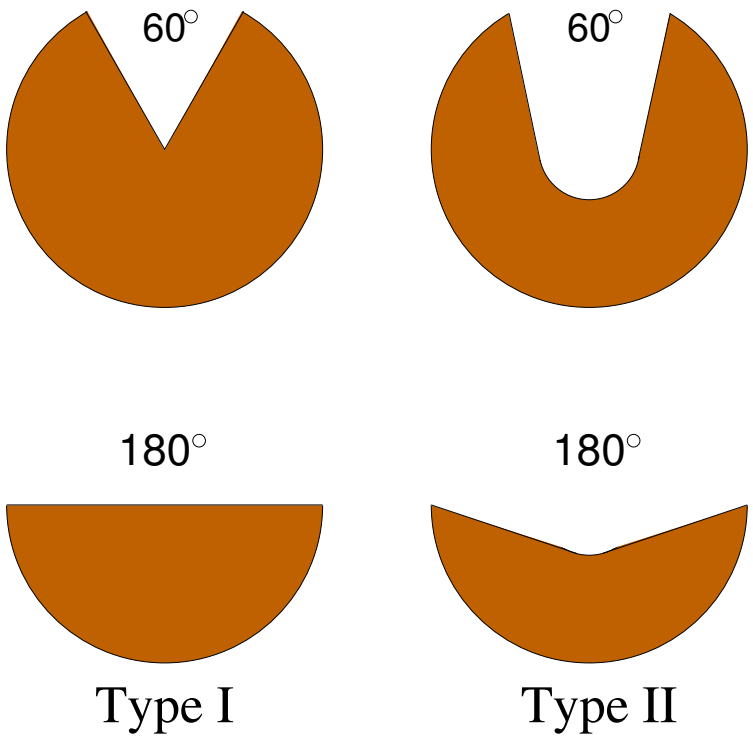

Figure 1.16: Pac-man shapes

pac-man shape is further divided into two types: type 1 (PM-I) is a circle with a wedge shaped slot removed while type 2 (PM-II) is a ring or torus with a slot removed [49–51]. The PM-I element was found to have a stable state with two distinct magnetic domains, which is undesirable for MTJs. However, the PM-II elements formed stable single domains regardless of slot angle, and the required switching fields were more consistent than for an array of

rectangular or hexagonal elements. A subsequent investigation of PM elements with a 40% larger aspect ratio determined that the vortex-driven domain reversal that was characteristic of the original PM shape would be replaced by a coherent reversal [52–54]. Coincidentally, an element resembling the elongated PM-II but referred to as a "rounded trapezoid" was also reported to have improved selectivity when compared to rectangular and ellipsoidal elements [37].

It is important to note that the simulations of various shapes described above assumed that the desired shape could be fabricated exactly as drawn, without significant defects or variation. Unfortunately, imperfections in the formation of the magnetic layers can have a profound effect on their characteristics. A cleft defect as small as 20 nm in the edge of a 700 nm square element was found to double the required switching field [55]. These edge defects arise from statistical variations in lithography and act as domain wall pinning sites. Of course, a square element is not representative of typical MTJs, but a similar study of circular nonmagnetic voids within elliptical shapes came to the same conclusion [56]. Voids with a diameter of 40 nm or 60 nm were created in an ellipse with axis lengths of 300 nm and 500 nm, and the magnetic switching was simulated for fields aligned at various angles to the easy axis. Voids close to the center of the ellipse had much less effect than those close to the edge, where a small void was found to double the required switching field for a field angle of 45°. Unfortunately, these studies evaluated only the micromagnetic behavior and did not determine the effect of imperfections on a tunneling junction.

1.6 Radiation Effects in Electronics

1.6.1 Total Ionizing Dose Effects

The vast majority of the particles encountered in space will be photons, electrons, and protons. They have sufficient kinetic energy to ionize a silicon atom by dislodging an

electron, but in general the quantity of free charge liberated will be small compared to the amount of charge that represents a distinct logic state in the circuit. These very light particles are so numerous that they are treated collectively and macroscopically as a bulk phenomenon, called *total ionizing dose* (TID).

Total ionizing dose (TID) effects are the result of accumulated damage in materials and change the characteristics of the material for long periods of time [57]. In digital integrated circuits the TID effects are seen as trapped charge in oxide layers because when an atom in an insulator is ionized the free electron will often be swept away by electric fields before recombination can occur, leaving behind a relatively fixed positive charge. Damage occurring at the interface between silicon and a gate oxide can also appear as a fixed negative charge by trapping holes.

To illustrate how TID changes the behavior of a MOS transistor, consider the top view of such a transistor shown in Figure 1.17. The light green background is the silicon substrate, which is

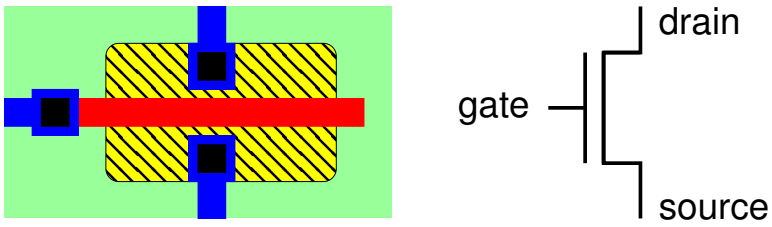

Figure 1.17: Top view of example transistor

commonly doped with a low concentration of acceptor impurities (a P- doping). The *field oxide*, a layer of silicon dioxide, SiO_2, is grown by oxidizing the silicon surface. The transistor is formed by first etching a hole in the field oxide down to the silicon surface, represented by the yellow area in the figure, and then regrowing a very thin *gate oxide*. The controlling structure of the transistor is its gate, shown in red, which is composed of doped polycrystalline silicon that is deposited on top of the field and gate oxides and lithographically patterned. At this point the transistor source and drain are formed by implanting donor atoms through the gate oxide. The field oxide and polysilicon gate block this implant, so only the

areas shown in yellow are affected. A donor impurity is used to implant the source and drain, and a sufficient concentration is used to counteract the light acceptor concentration in the silicon and create a thin heavily doped (N+) layer. Now connections can be made to the transistor. The transistor is covered with a new dielectric layer, and holes in this layer (referred to as *contacts* and shown as black squares) are etched where connections to the underlying transistor will be made. The holes are typically filled with a metal such as titanium or tungsten, and then connecting wires (in blue) are patterned using aluminum or copper alloys.

Without any voltage applied to the gate there can be no current flow between the source and drain, because depletion regions exist at the PN junction formed where the N+ regions meet the P- substrate. The substrate area under the polysilicon gate is also P- since the gate blocked the source/drain implant. The substrate will normally be connected to the most negative voltage in the circuit to insure that these PN junctions are never forward biased, a technique known as junction isolation. To "turn on" the transistor and allow current flow between the source and drain a positive voltage is applied to the gate. A large electric field is created across the thin gate oxide, and electrons will be attracted to the silicon surface under the gate. If the voltage on the gate is large enough, so many electrons are pulled to the surface that they overcome the acceptor impurities and *invert* the surface. The inversion layer now behaves as if it were also N+ and forms a conducting channel between the source and drain, so this structure is called an "n-channel" field-effect transistor. Originally, such transistors used aluminum instead of polysilicon for the gate material and were called n-channel metal-oxide-semiconductor, or NMOS, transistors. Although the vast majority of field-effect transistors currently manufactured do not use metal gates, the "MOS" reference continues to be applied to field-effect transistors in general. By exchanging the donor and acceptor impurities a p-channel (PMOS) transistor can be created, which will be turned on by applying a relatively negative voltage to the gate. The voltage needed to turn on the transistor is its *threshold voltage*, specified with respect to the source terminal. NMOS transistors have a positive threshold voltage and PMOS transistors have a negative threshold voltage, but the

magnitude in both cases is typically 0.5 V or less.

When a light energetic particle passes through the silicon layers of the substrate or source/drain regions a few atoms are ionized but this is insignificant when compared to the concentration of free carriers contributed by the donor or acceptor impurities. However, when ionization occurs within the field or gate oxide the characteristics of the transistor may be permanently altered, as shown in Figure 1.18. This figure shows a cross section of the

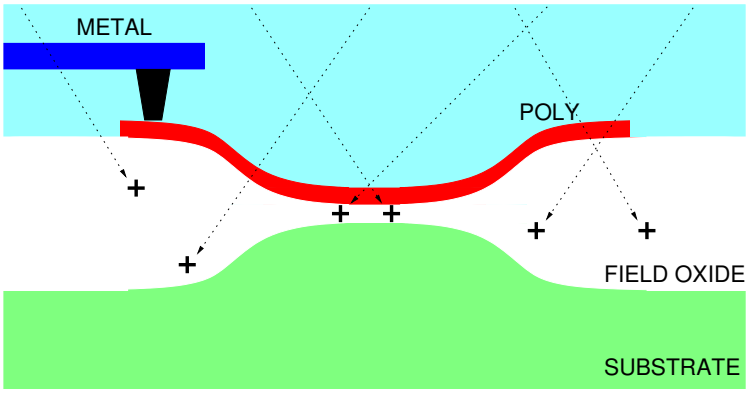

Figure 1.18: Charge accumulation due to ionizing radiation

transistor in Figure 1.17, taken along the long dimension of the polysilicon gate. If the ionized atoms and free electrons appear within the oxide and there is no electric field present then they are most likely to recombine, causing no harm. However, if the ionization occurs in the presence of an electric field the highly mobile electrons will be swept away before recombination can occur, leaving behind the ionized atom as a trapped positive charge. As trapped charge accumulates it acts much like a positive voltage applied to the transistor's gate: it tends to turn on NMOS transistors but make it more difficult to turn on the PMOS transistor. In other words, the trapped charge makes the transistor thresholds more negative.

A secondary TID effect is the result of damage at the interface between the gate oxide and the silicon substrate. Fixed interface states may accumulate at the surface, resulting in a decrease in carrier mobility through the channel and a positive shift in transistor thresholds [58]. A number of factors have been identified that influence the magnitude of total dose effects that

any particular integrated circuit will exhibit, including seemingly unrelated issues such as plastic packaging and burn-in [57].

Fortunately, as integrated circuit technology inexorably moves to smaller feature sizes the TID damage becomes less serious [59, 60]. Very thin gate oxides provide a smaller volume for accumulated positive charge, and the improved quality of the interface between the oxide and the substrate results in fewer interface states. For most civilian spacecraft missions using modern CMOS technologies the TID effects will not be a critical factor.

Floating gate memory technologies, such as Flash, can be seriously degraded by TID [61]. Information is stored by using tunneling through thin insulators to deposit charge on an unconnected, or "floating", gate electrode, thereby altering the threshold voltage of the associated transistor. The electrons and holes created by the ionizing radiation tend to neutralize whatever charge was left on the floating gate during programming, rendering it impossible to successfully discern the '1' state from the '0' state. Failures have been seen with TID levels of around 1 kGy (100 krad) in commercial Flash memories. Unfortunately, scaling to thinner oxides in advanced devices does not necessarily eliminate this problem since the amount of charge representing the stored data will be reduced to tens or hundreds of electrons [62, 63].

1.6.2 Single Event Effects

In addition to numerous electrons and protons the space environment also contains heavier atomic nuclei with sufficient kinetic energy to easily penetrate the metal skin of a spacecraft, and these are commonly referred to as cosmic rays or cosmic particles. Cosmic rays are encountered infrequently, with a probability that decreases rapidly for heavier nuclei. In general, the effects of a cosmic particle can be treated as though they are localized in both space and time and are consequently called *single event effects*.

Single event effects (SEE) include a variety of transient and permanent failures, with the

common thread that they are caused by a single, highly energetic particle that passes through a circuit. As the particle travels through a material it can ionize atoms, creating a trail of charge. The amount of charge left behind depends on the mass and energy of the particle, as well as the nature of the material through which it passes. Years ago, the release of alpha particles (helium nuclei) from naturally occurring radioactive decay in packaging materials resulted in transient failures in commercial integrated circuits. Today, the SEE threat is most important in spacecraft but terrestrial single events may become significant again as transistor feature sizes approach the nanometer regime.

The Cosmic Particle Environment of Space

The ability of a cosmic particle to deposit charge by ionizing atoms as it passes through some material is called its stopping power or linear energy transfer (LET) [64]. The LET of a particle is measured in $\mathrm{MeV \cdot cm^2/mg}$, which can then be multiplied by the density of the target material to determine how much energy will be transferred to that material per unit length of the particle track. For example, a particle with an LET of $1.0\ \mathrm{MeV \cdot cm^2/mg}$ passing through silicon ($2.33\ \mathrm{g/cm^3}$) will leave behind $2330.0\ \mathrm{MeV/cm}$ or $0.233\ \mathrm{MeV/\mu m}$. In silicon 3.6 eV of energy is needed to dislodge an electron, so the example particle will create 6.47×10^4 electron-hole pairs per μm of track length. The equivalent charge of an electron is 1.6×10^{-19} C, so the $1.0\ \mathrm{MeV \cdot cm^2/mg}$ particle deposits $10.4\ \mathrm{fC/\mu m}$.

In the space environment, the SEE threat varies significantly with the altitude and orientation of orbit, or with the path of a deep space mission. The most common particles of interest are protons and electrons that have been trapped and accelerated in the earth's magnetic field, creating the Van Allen radiation belts. Of these, protons have sufficient momentum to cause SEE in electronics. The outer Van Allen belt also traps significant numbers of heavier particles called galactic cosmic rays (GCR) which are the nuclei of naturally occurring elements. At large distances from earth the primary threat is from GCR, although radiation

belts can be found around other bodies that have a magnetic field. A typical cosmic particle

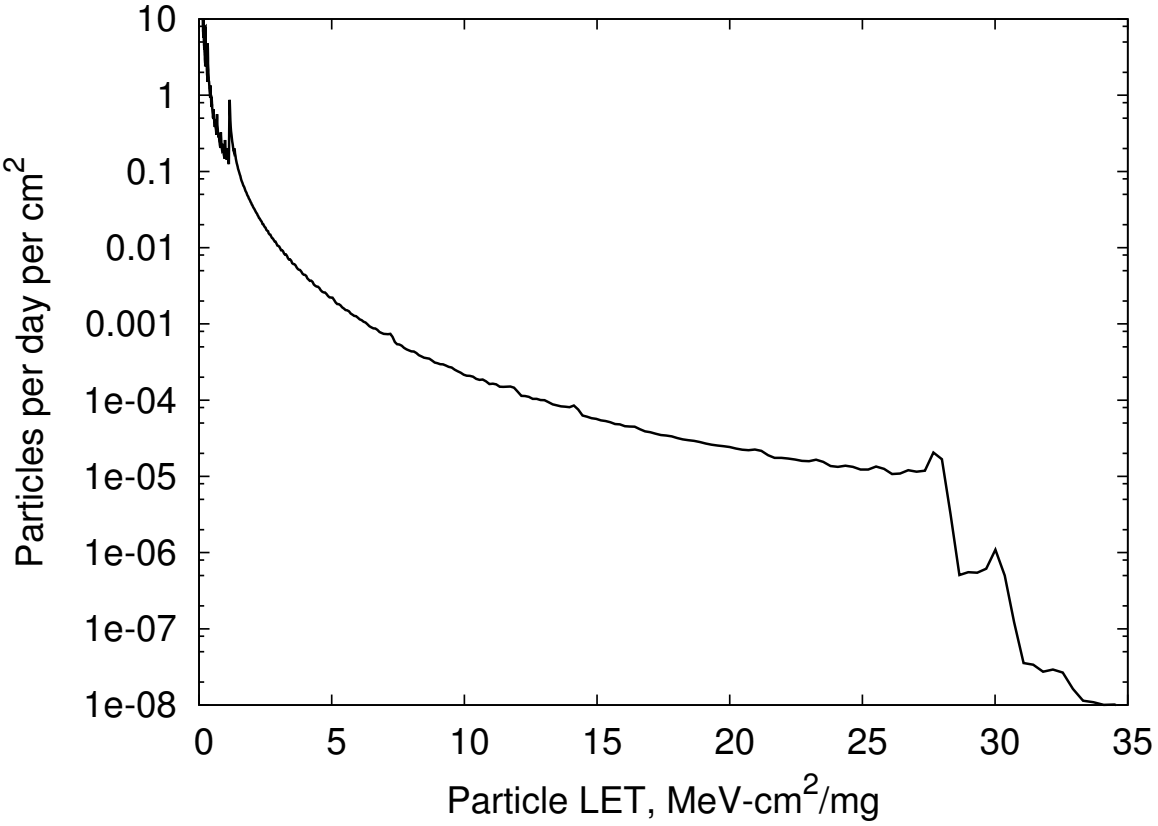

Figure 1.19: Example cosmic particle spectrum

spectrum is shown in Figure 1.19 for a spacecraft in geosynchronous orbit during a period of minimal solar activity [65]. Note that the likelihood of encountering a particle falls dramatically as the LET increases.

The non-uniformities in the cosmic particle spectrum resemble the relative abundance of the elements in the universe as a function of their atomic number, which is in turn a reflection of the various processes involved in the evolution of stars. The elements up to iron and nickel are created by the normal gravity-driven fusion process, but heavier elements require a supernova explosion and are much less common in space [66, 67]. Figure 1.20 shows that 99.8% of the atoms in the universe are hydrogen or helium, and less than 1% of the atoms are heavier than nickel (Z=28). Note that the light elements lithium, beryllium, and boron (Z numbers 3, 4, and 5) are consumed quickly, and at relatively low temperature, to create carbon and are therefore

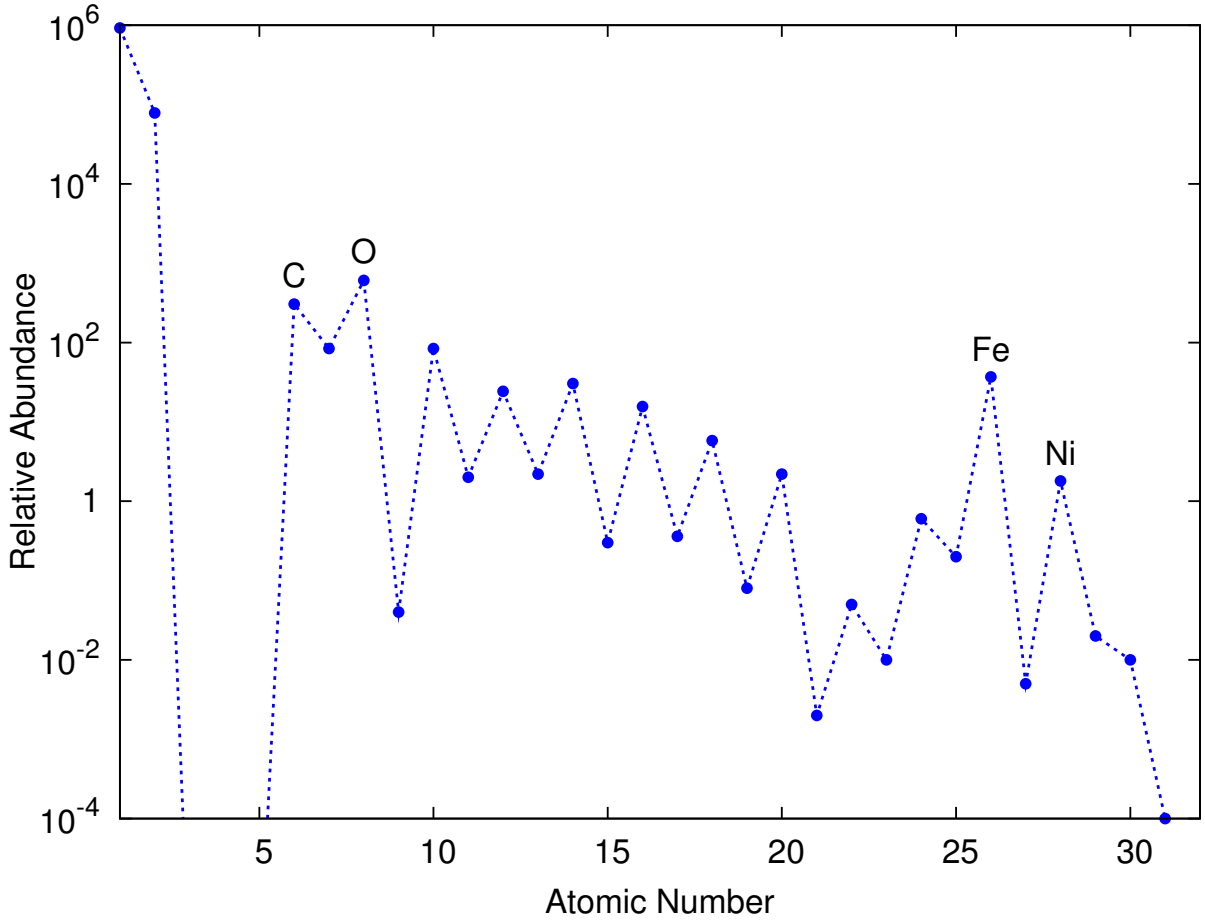

Figure 1.20: Relative abundance of the elements, per 10^6 atoms [67]

scarce. The elements in the "iron peak" region account for only about 15 of every million atoms, and beyond this peak the heavier elements are quite scarce. For practical purposes the maximum cosmic particle LET encountered in space is 28.68 $\text{MeV}\cdot\text{cm}^2/\text{mg}$ which is delivered by a 79.15 MeV ${}_{26}\text{Fe}^{56}$ nucleus [68]. For comparison, the peak LET of an alpha particle (i.e. ${}_{2}\text{He}^{4}$ nucleus) is 1.57 $\text{MeV}\cdot\text{cm}^2/\text{mg}$ and for a ${}_{92}\text{U}^{238}$ nucleus the peak LET is 105.54 $\text{MeV}\cdot\text{cm}^2/\text{mg}$.

Single Event Charge Collection in Bulk CMOS

The total amount of charge deposited is obviously proportional to the track length over which charge can be collected without recombination. For design purposes the particle flux is

assumed to be isotropic, so it is possible that a particle could pass through a depletion region along its longest dimension. Furthermore, the deposited charge can distort the depletion region in what is known as *funneling* or an *ion shunt* and the effective track length may be significantly longer than the original depletion region [69–73]. An effective path length of several microns is not unlikely and a conservative design approach might assume a maximum track length of 10 µm, which when multiplied by a conservative maximum LET estimate of 28.68 $\mathrm{MeV \cdot cm^2/mg}$ yields a maximum charge deposition of approximately 3.0 pC in silicon [74, 75].

The impact of this deposited charge is largely determined by where the particle strike occurs, as shown in Figure 1.21. If the ionization trail is created within an equipotential region of bulk

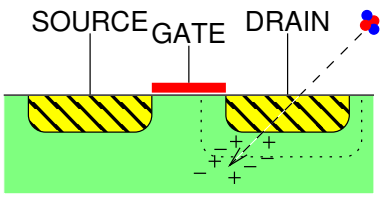

Figure 1.21: Cosmic particle striking a transistor

silicon then the hole/electron pairs will tend to recombine without producing any observable effect. However, if the ionization occurs within the depletion region of a reverse-biased PN junction, as shown in the figure, then the electric field across the depletion region will quickly sweep the carriers apart and little recombination can take place. The net effect is the injection of charge into a reverse-biased transistor drain.

There are two important macroscopic observations regarding a cosmic particle strike. First, the most sensitive areas of a circuit are the reverse-biased source or drain regions, which typically will be the NMOS sources and drains at the supply voltage and PMOS sources and drains at ground. Second, the freed carriers will flow in direction dictated by the applied electric field and will therefore tend to forward bias the struck junction. In other words, PMOS drains will be most vulnerable when at the ground potential and the particle strike will

cause the node voltage to become more positive. Similarly, NMOS drains are most vulnerable when at a high potential and will be driven toward ground.

Single Event Effects in SOI Transistors

Although *bulk* CMOS, where transistors are formed by implanting dopants into a thick conducting substrate, has constituted the lion's share of the market, silicon-on-insulator (SOI) technologies offer several advantages. The cost of manufacturing SOI has been reduced to the extent that SOI is a viable contender for high-volume commercial products, which means that aerospace applications will also benefit from the economy of scale. A typical SOI transistor cross-section and equivalent circuit are shown in Fig. 1.22. As the name suggests, SOI

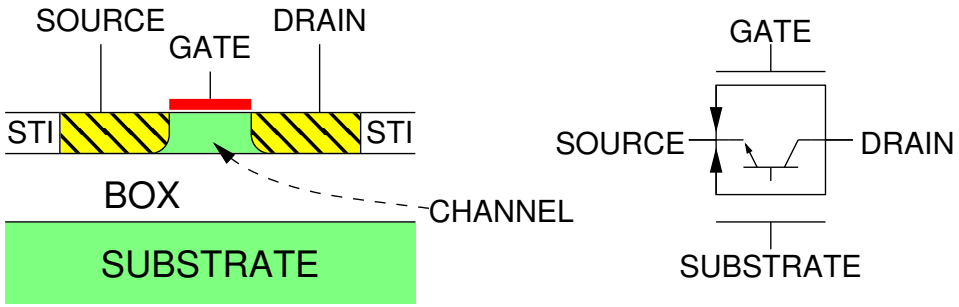

Figure 1.22: Generic SOI transistor

transistors are manufactured on top of an oxide layer, which is often created by deeply implanting oxygen into a silicon wafer and is therefore called a buried oxide (BOX). The BOX is typically < 500 nm thick, and is created up to 200 nm below the silicon surface [76]. Transistors are formed in the silicon surface layer, and isolated from one another laterally with a shallow trench isolation (STI) oxide.

The equivalent circuit of the SOI transistor is shown on the right side of Fig. 1.22. In addition to the intended MOSFET, there is a parasitic MOSFET formed at the bottom of the channel, with the BOX as the gate oxide and the substrate as the gate terminal. This back-gate transistor will typically have a relatively high threshold voltage and can be neglected under

normal operating conditions. There is also a parasitic bipolar transistor, with the MOSFET source and drain as its emitter and collector and the MOSFET channel as its base. As the MOSFET channel length decreases the existence of the parasitic bipolar becomes more significant, and it is the primary source of the single-event effects in SOI.

Because most of the depletion regions that exist in bulk CMOS have been eliminated in SOI, the amount of charge that is collected by a cosmic particle strike is greatly reduced [77]. However, if charge is deposited in the channel it may forward bias the channel-source (base-emitter) junction and cause bipolar amplification of the current. This effect is most severe in transistors built in a relatively thick silicon layer (> 100 nm), which is known as partially depleted (PD) SOI because the depletion region in the channel does not extend completely down to the BOX. The depletion region isolates a thin layer of the channel near the BOX, which is often called a "floating body" and introduces several undesirable effects. The bipolar transistor existing in the floating body has a relatively high current gain ($\beta \approx 100$) and significantly amplifies any charge collected from a cosmic particle. An obvious technique for disabling the bipolar transistor is to add body ties that hold the channel potential very near the source potential. However, the manufacturing processes used to create the MOSFETs often make it difficult or impossible to add body ties. When ties are used they are often placed at one or both ends of the channel, and the large resistance that exists from the ties to the center of the channel greatly reduces their effectiveness.

Thinning the silicon layer to less than 100 nm typically results in fully depleted (FD) SOI, and eliminates the floating body. The parasitic bipolar transistor still exists, but its gain is greatly reduced and less charge can be deposited in the thinner channel layer [78]. However, current amplification is still possible [79] and leads to significant single event effects.

Single Event Transients and Single Event Upset

To illustrate the single event effects (SEE) that occur in circuits as a result of cosmic particle strikes, consider the circuit shown in Figure 1.23. The top portion of the figure shows a cross-section of the silicon layers making up the circuit. Since this circuit uses both NMOS and PMOS transistors it is known as a *complementary* MOS, or CMOS, circuit. As before, fabrication begins with a lightly doped P- substrate, but for CMOS it is necessary to implant a deep, lightly doped N- region called the *n-well* where the PMOS transistors will be formed. The NMOS and PMOS transistors are then created in separate sequences of lithography, masking, and impurity doping. The connections between the transistors are shown schematically above the cross section view, although it is implied that the substrate will be connected to V_{DD}, the most positive power supply potential in the circuit, and the n-well will be connected to V_{SS}, the most negative potential (usually ground). Below the cross section

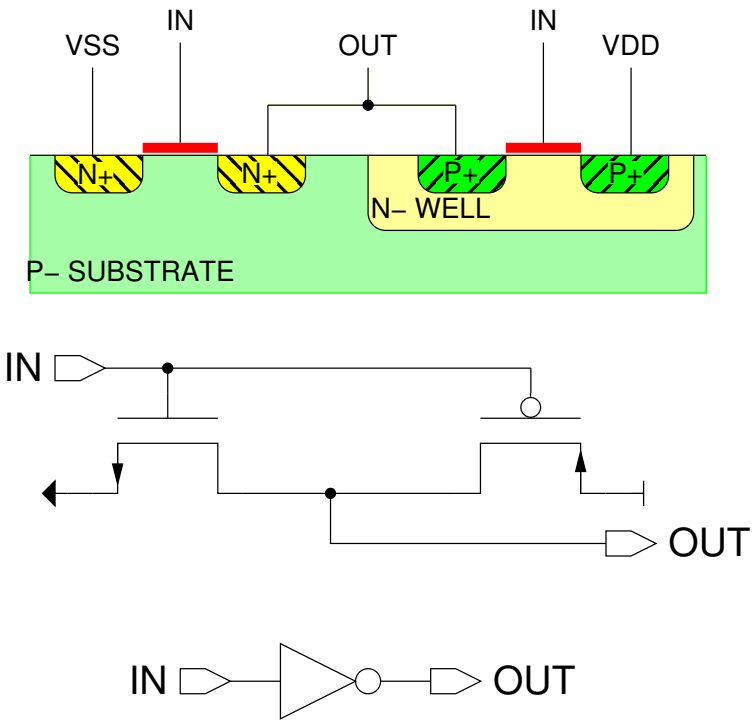

Figure 1.23: CMOS inverter

view in the figure is the equivalent schematic representation for the two transistors, and at the

bottom of the figure is the conventional schematic symbol for this widely used circuit. When the node labeled IN is at the V_{DD} potential the NMOS transistor will be turned on (conducting) while the PMOS transistor is turned off, so the OUT node will be pulled to V_{SS} through the NMOS transistor. Similarly, when IN is at V_{SS} the NMOS transistor is off and the PMOS is on, pulling OUT to V_{DD} through the PMOS transistor. If the V_{DD} potential is interpreted as a Boolean '1' and V_{SS} is interpreted as a '0' then this circuit performs a logical inversion and it is therefore called an *inverter*.

Figure 1.24 shows a series connection of two inverters, where the input of the first inverter is connected directly to V_{DD}. In this example V_{DD} is 2.5 V while V_{SS} is connected to ground. The waveforms in the figure illustrate how the inverters respond to a simulated cosmic particle strike at node A, the output of the first inverter. The strike occurs at 1 ns in the simulation and the deposited charge quickly drives node A more positive than V_{DD}. The output of the second inverter, node B, falls to V_{SS} when it inverts the logical value at its input. The pulse on B is an indirect consequence of the cosmic particle known as a single event transient (SET). Since the input of the first inverter remains at V_{DD} its NMOS transistor remains turned on and it eventually dissipates the deposited charge, returning node A to its correct state. Shortly thereafter node B also returns to its desired state, and the single event has ended.

A common method for creating a memory element is to cross-couple two inverters, connecting the output of each to the input of the other. Once forced into a given logic state the positive feedback of this arrangement maintains that logic state indefinitely. Figure 1.25 shows a simulated cosmic particle strike to such a latch circuit. As before, node A is forced high causing node B to fall to ground. Since B is now connected back to the input of the first inverter, its PMOS transistor is turned on rather than its NMOS transistor and there is no current path to restore the original state of the circuit. Instead, the incorrect state is maintained by the positive feedback and becomes the stable state of the latch. When a cosmic particle strike changes the stored value in a memory element a single event upset (SEU) is said to

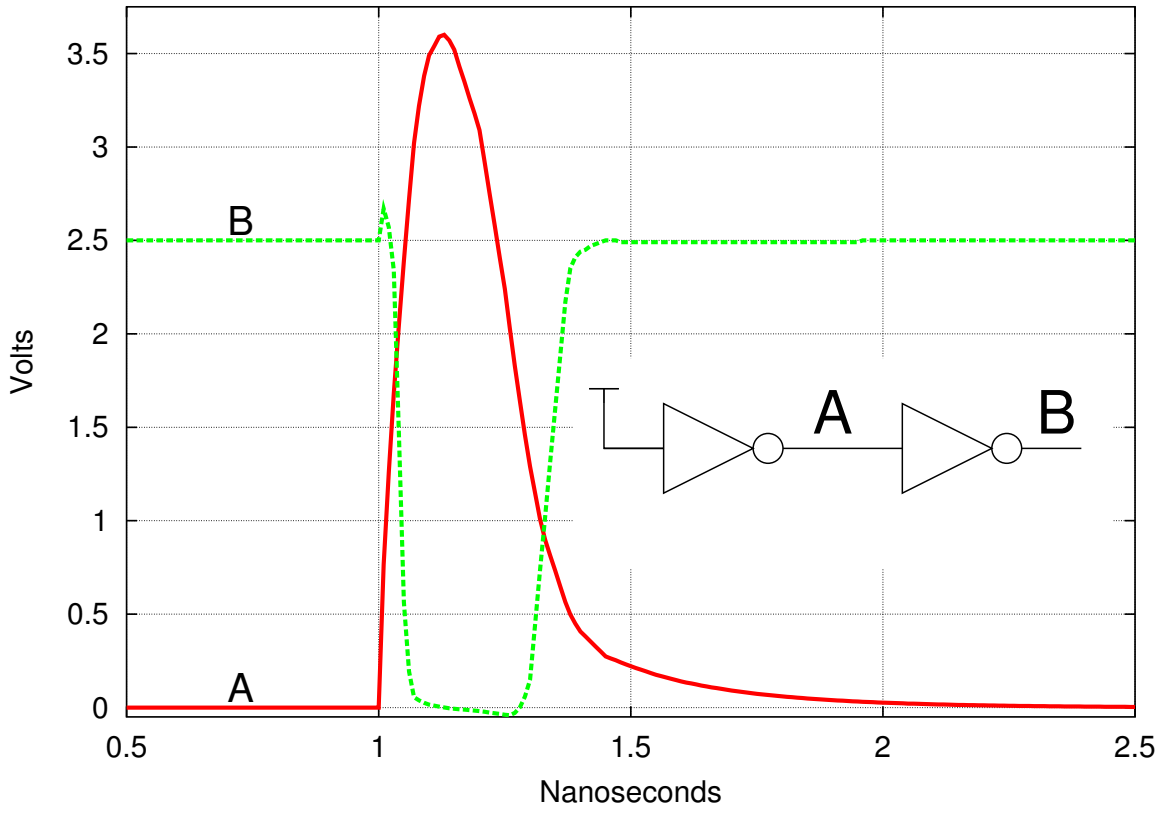

Figure 1.24: Cosmic particle strike to inverter pair

occur.

Obviously, an SEU in the critical circuitry of a spacecraft could have disastrous consequences. Rocket engines might fire at the incorrect time or fail to fire when necessary. Communications with the Earth could be disabled when a transmitter or receiver behaves incorrectly. Although the possibility of SEU was suggested in 1962 it was not until 1975 that such a failure was observed in a communications satellite [64]. Since that time considerable effort has been devoted to the development of SEU-resistant and SEU-immune circuits, and these are discussed at length in Section 1.6.3.

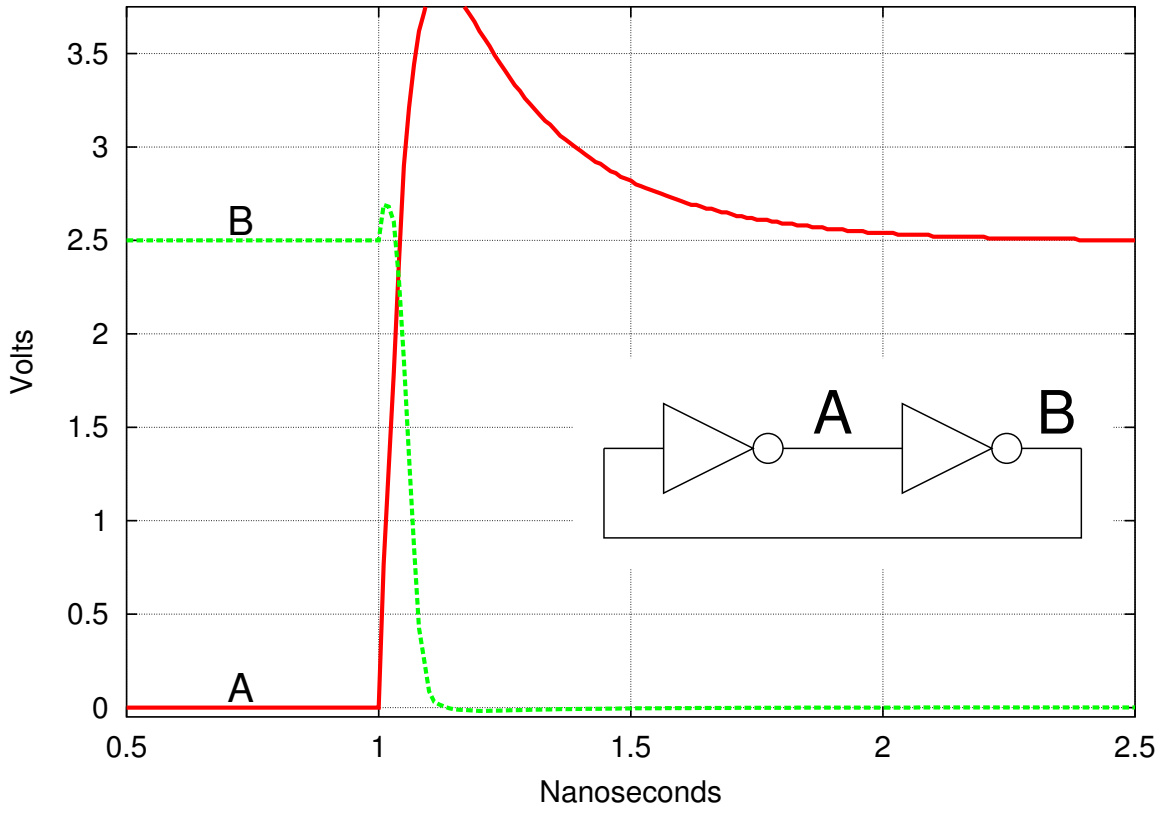

Figure 1.25: Cosmic particle strike to latch

Single Event Latchup and Snapback

A common, and very serious, single event effect in commercial CMOS products is SEL, or single event latchup. This occurs when the cosmic particle deposits sufficient charge within the transistor structure to trigger the parasitic bipolar transistors into a self-sustaining high current state [80]. The physical location and equivalent schematic of these bipolar transistors is shown in Fig. 1.26, which has been simplified to show only those structures that contribute to latchup. Starting from the left side of the figure are two implants into the lightly doped P-type substrate: a P+ ohmic substrate contact and an N+ FET source region. An N-type well contains a P+ FET source region and an N+ ohmic well contact. The substrate contact and well contact serve to maintain a reverse-biased junction between the substrate and well. The transistor source regions are typically packed tightly together, while the substrate and well

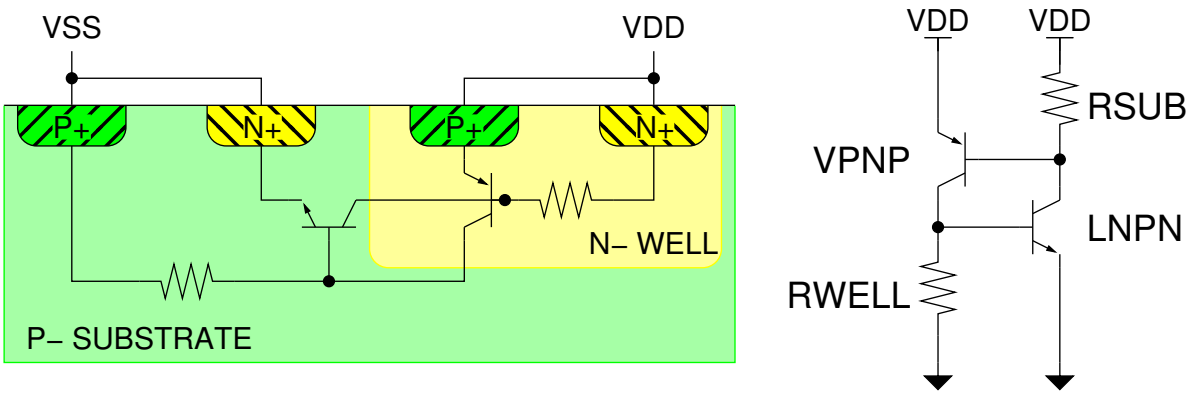

Figure 1.26: Parasitic bipolar transistors in CMOS

contacts may be tens of micrometers away. Parasitic bipolar transistors form around the sources, with a lateral NPN from the N+ source to the N- well and a vertical PNP from the P+ source to the P- substrate. Because the substrate and well are more lightly doped, and the contacts are located some distance away, there is a significant resistance from them to the base terminals of the parasitic bipolar transistors.

The equivalent circuit is given on the right side of Fig. 1.26. Without the resistors this circuit resembles the PNPN structure known as a silicon controlled rectifier, or SCR. As long as the voltage drop across both resistors is insufficient to forward bias the base-emitter junction of a transistor, no current will flow. SEL can occur when a cosmic particle injects charge into the substrate or well that results in a large current flowing through the respective resistor. Once either of the bipolar transistors starts to conduct a significant collector current it will produce a voltage across the other parasitic resistor, causing collector current to flow through the other transistor, which then reinforces the voltage drop across the resistor that was struck by the cosmic particle. This condition will be self-sustaining if the product of the current gains of the two transistors is greater than unity ($\beta_N \beta_P > 1$) and the supply voltage is somewhat greater than the sum of the base-emitter voltages ($V_{DD} > V_{BEN} + V_{BEP}$). The combined collector current during latchup can be fairly large, causing local voltage drops, spot heating, or even permanent physical damage. Once entered, the latchup state can be terminated only by lowering the supply voltage, which is inconvenient in practice.

Fortunately, there are simple and effective strategies for preventing SEL. In bulk CMOS, using heavily doped substrates with thin epitaxial layers can greatly reduce the substrate resistance [81], and the judicious placement of substrate/well contacts or guard rings is a common cure [82].

Using SOI technology clearly eliminates any threat of latchup since the parasitic four-layer structure no longer exists. However, a similar phenomenon known as single event snapback (SES) has been observed in SOI transistors [77]. If a very high electric field exists near the drain junction, impact ionization can result in additional charge injection into the base, which is regenerative and can only be stopped by lowering the drain voltage. Fortunately, the drain voltage required to sustain snapback is relatively high and SES should not be a concern in typical circuits.

Permanent Damage Effects

It is also possible for a single heavy particle to permanently damage some structures in integrated circuits, although these effects are not usually a threat to CMOS logic circuits. Single heavy ions have been observed to cause permanent damage to Flash memories [83]. As with TID, floating gate memories are susceptible to data loss caused by SEE. One proposed model for these failures is that by ionizing atoms along its path the particle creates a conductive “pipe” between the floating gate and the control gate of the transistor, thus draining the stored charge [84]. Measurable charge loss can result from the strike of a single alpha particle, so this is a reliability concern for almost any space application. Again, it appears that this effect will only become more pronounced as device geometries shrink and the amount of charge stored on the floating gate is reduced. Permanent changes have also been observed in floating gate cells, where the retention characteristics of the cell are significantly degraded [62].

It is also well known that heavy ions, such as cosmic rays, can cause permanent and

significant damage to thin insulating films. If a large electric field is applied across the insulator and a heavy ion passes through it, a catastrophic failure known as single event gate rupture (SEGR) can occur [85]. This is a common failure mode for high-voltage MOS transistors in space, but is becoming increasing relevant as CMOS technology scales to smaller features and thinner oxides [86, 87]. SEGR is most likely to occur for cosmic particles with high LET and with electric fields above 1 MV/cm, which is consistent with the electric fields applied when reading the state of an MTJ. Unfortunately, research in this area has focused on insulators such as SiO_2 and SiN rather than Al_2O_3 or MgO so the significance of SEGR in tunneling junctions is yet to be determined.

1.6.3 SEU-Resistant Latch Circuits

There is considerable variation in the radiation tolerance of commercial integrated circuits so the first step in hardening a spacecraft's electronics is selecting those components that offer the greatest inherent resistance. For example, dynamic RAM stores information as a very small charge on a capacitor and is much more sensitive to SEU than static RAM, which has active transistor drive for storing data.

Unfortunately, as commercial integrated circuit manufacturers advance their products to higher levels of miniaturization it becomes increasingly difficult to identify commercial products that are suitable for use in space. Submicron transistors and wires have very small nodal capacitance, and reliability concerns force the use of lower and lower supply voltages, with the net result that the amount of charge used to represent a logic state is dwarfed by the charge injected by a cosmic particle. In order to deploy modern electronic systems in spacecraft many of their components must be specifically modified to increase their SEU resistance.

One of the earliest, and most successful, techniques for hardening CMOS latches is *resistor hardening* [88]. As shown in Figure 1.27 a resistor is placed in the feedback path of the latch

circuit, creating an additional node, C, in the circuit where information is stored. The essence

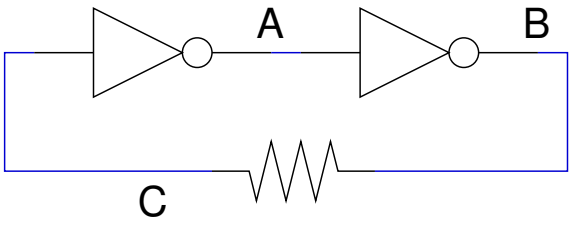

Figure 1.27: Latch with resistor hardening

of this approach is that the added resistor, along with the input capacitance of the left inverter, forms a low-pass filter. Voltage transients occurring at nodes A and B as a result of cosmic particles will not change the logic level stored at node C unless their duration exceeds the RC time constant of the filter. By simulating a worst-case particle strike at nodes A and B the maximum recovery time of these nodes can be determined, and the value of the resistor is chosen to assure that the filter time constant is greater than the recovery time.

For resistor hardening to be effective there must be no cosmic particle charge collection at C, which implies that the resistor can not be formed using a junction-isolated structure in the silicon substrate. Typically, these resistors will be made using lightly doped polysilicon, either by masking the doping implants to the desired portions of the same polysilicon layer that forms the transistor gates or by adding a second polysilicon layer that is lightly doped everywhere. When the necessary resistance exceeds 100 kΩ it becomes difficult to manufacture polysilicon resistors that are both small and consistent in value. Furthermore, these lightly doped resistors have a large negative temperature coefficient [89, 90]. Consequently, the target resistor value is typically much larger than the minimum value needed for SEU resistance.

A significant disadvantage of resistor hardening is that the amount of time required to write data into the latch increases along with the RC time constant of the filter. Although the cosmic particle passes through the circuit and creates free hole/electron pairs in a few tens of picoseconds, the process of collecting and removing the deposited charge may continue for

nearly a nanosecond and the filter must protect node C until this process is complete [91]. At some point the write time penalty becomes too severe, and resistor hardening loses its appeal [92].

An alternative to resistor hardening that is becoming increasingly popular is known as *design hardening*. The advantage of design hardening is that it does not require special process steps or capabilities, such as a lightly-doped polysilicon layer, to achieve a high level of resistance to radiation effects [93, 94]. Instead, unique circuit configurations have been developed to add information redundancy to memory cells.

The DICE (Dual Interlocked CEll) circuit is one of the most widely used design-hardened latch configurations, in part because its use is not encumbered by patents [95–97]. As shown in Fig. 1.28, the cell comprises four storage nodes, each connected to an NMOS transistor and

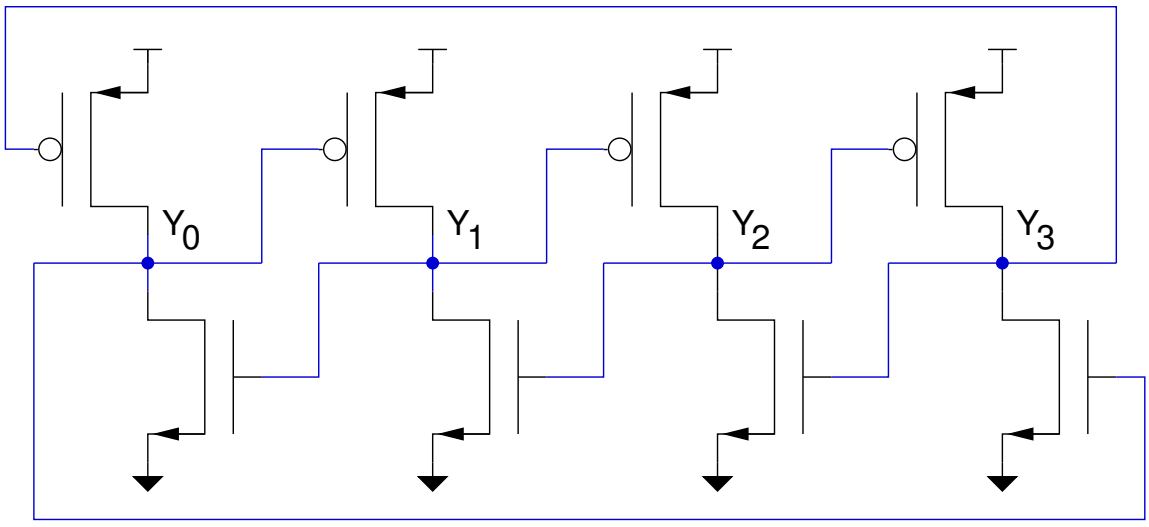

Figure 1.28: DICE schematic

a PMOS transistor. There are two stable states for this cell:

$$Y_0Y_1Y_2Y_3 \in \{1010, 0101\} \tag{1.30}$$

A critical design criterion for the DICE is that one type of transistor, either the NMOS or the PMOS, must have significantly more drive than the other under all conditions. In general, this

condition is most conveniently satisfied by assuming that the NMOS is stronger than the PMOS, which is typically the case for similar transistor sizes due to the higher mobility of electrons. This simplified schematic does not include the circuitry needed to load new values into the latch, which is typically accomplished using pass transistors that can overdrive either the NMOS or PMOS transistors in the heart of the latch itself.

To see how the cell recovers from a cosmic particle strike, consider the case where $Y_0Y_1Y_2Y_3 = 1010$ and a particle strikes Y_1, changing it from a 0 to a 1. The PMOS pulling Y_2 high is turned off, leaving Y_2 floating at a (correct) high level due to the node capacitance. The single event at Y_2 also turns on the NMOS with its drain at Y_0, causing contention between the PMOS and NMOS at Y_0. Assuming that the NMOS is stronger than the PMOS, Y_0 will be pulled to an incorrect low level. The low level at Y_0 disables the NMOS pulling Y_3 low, so Y_3 floats but maintains a correct low level due to node capacitance. The low level on Y_0 also enables the PMOS that pulls Y_1 high, but Y_1 soon returns to the desired low level because the NMOS pulling this node low is enabled and is stronger than the PMOS. As the correct state is restored on Y_1 the NMOS driving Y_0 will be disabled; Y_0 then also returns to its correct state. The state transitions in this case were

$$1010 \underset{hit}{\longrightarrow} 1110 \longrightarrow 0110 \underset{recover}{\longrightarrow} 0010 \longrightarrow 1010 \tag{1.31}$$

Now consider $Y_0Y_1Y_2Y_3 = 0101$ and a particle striking Y_1 causes it to change from a 1 to a 0. The PMOS driving Y_2 is enabled but node Y_2 does not change state because the PMOS is weaker than the NMOS. The NMOS transistor driving Y_0 is disabled but Y_0 floats at its correct value. Since neither Y_0 nor Y_2 have changed to an incorrect state, the desired level on Y_1 is soon restored. Since only Y_1 enters an incorrect state the state transitions are simply

$$0101 \underset{hit}{\longrightarrow} 0001 \underset{recover}{\longrightarrow} 0101 \tag{1.32}$$

The SERT (Single Event Resistant Topology) cell shown in Fig. 1.29 was designed to overcome two of the disadvantages of the DICE. By adding an additional transistor to each of

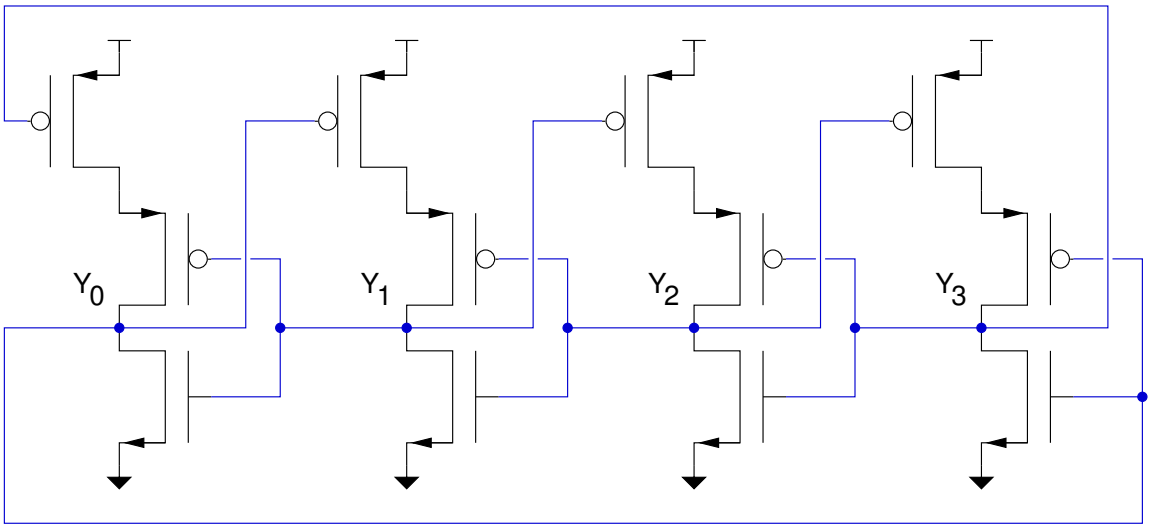

Figure 1.29: SERT cell schematic

the four sections of the cell, there is no need to insure that one transistor type is always much stronger than the other. This is important in situations where the relative strength of the transistors may be dynamically adjusted, or if the leakage current of the transistors is significant when compared to the saturated drive current. Because the SERT cell never experiences contention, where a node is simultaneously pulled high and low, it recovers slightly faster than the DICE [98, 99]. While Fig. 1.29 illustrates a SERT cell with an additional PMOS in each section with its gate connected to the gate of the DICE NMOS, an alternative configuration adds an NMOS transistor instead, with its gate tied to the gate of the DICE PMOS. Of course, it is possible to add both a PMOS and an NMOS to the original DICE configuration and this circuit has been patented. . . twice [100, 101].

Single event upsets can also occur in SOI latches and RAM cells. They are more severe in partially-depleted SOI but also occur when the channel is fully depleted. Therefore, many of the SEU-resistant circuit designs discussed above are applicable to SOI circuits, and a few configurations that are unique to SOI have also been developed [102].

Unfortunately, as integrated circuit dimensions shrink and transistors are packed ever closer to each other, one of the most fundamental assumptions underlying these latch designs fails to hold and it becomes possible for a single particle to deposit significant charge on more than one circuit node. The brute-force solution to this problem is to identify those critical node pairs, where a small charge collected at both nodes would be more serious than a larger charge on either node alone, and physically separate them in the layout. Simulations performed at the process level for feature sizes from 0.5 μm down to 90 nm suggested that "dummy" source/drain junctions placed near a sensitive node could significantly reduce the collected charge, allowing more efficient layout [103]. Dummy drains that were reverse biased with respect to the underlying material provided superior protection to those that were not reverse biased. A more extensive study evaluated the improvement obtained by extending the substrate/well contacts to create "guard contacts" between transistors as well as partitioning the N-well to isolate the PMOS devices. These simulations, as well as actual heavy-ion testing, found that charge collection in PMOS transistors is significantly reduced by adding N+ guard contacts between the transistors in the N- well, but the guard contacts between NMOS transistors and the partitioned well provided little improvement [104]. Unlike the dummy source/drain structures, the guard contacts were of the same type (P or N) as the underlying substrate or well and therefore could not be reverse-biased.

Chapter 2

Design of a Radiation-Tolerant MTJ Memory Cell

The design process is always driven by opposing forces. On one hand there are the desires and needs of the end user, which are often completely independent of any technical issues. On the other hand, the laws of nature place limits on what is achievable, or at the very least they present trade-offs that must be evaluated. Both of these influences are explored in this chapter, in the context of the current research.

Other researchers have faced similar issues, so it is instructive to examine the solutions they developed. In the last section of this chapter the capabilities of several examples of the prior art will be compared to the goals for this research. This discussion will introduce, by example, the criteria and metrics that will guide the development of new circuits.

2.1 Design Goals

The goal of this work is to develop a magnetic memory cell, or latch, which can be used in a generic digital design to replace replace conventional latches in a circuit when non-volatile storage is needed. Unlike an MRAM, these latches will be interspersed with logic functions. They will be compatible with conventional design tools for integrated circuit synthesis and physical design.

2.1.1 Functionality

The embedded memory cell should have at least one signal output that continuously provides either the true value or complement value of the data value stored in the latch. If the true value is provided, this signal would typically be called the Q output of the latch. If instead the complement, or inverted, data value is provided then the signal might be called QN (for Q-not), QB (for Q-bar), or $\overline{Q}$; the designation of QN will be used here. The memory cell must provide one of these outputs, but it may optionally provide both.

There must be a mechanism for storing a data value in one or more MTJs in the memory cell, so that the stored data is retained if power is removed from the cell. The cell must have an input signal that specifies the value to be stored, and one or more control signal inputs that cause the value to be stored. To reduce the number of signal lines in practice it is advantageous if the control signals can be connected to a number of memory cells in parallel, rather than requiring a distinct set of control signals for each latch.

Finally, there must be a mechanism for transferring the data stored in the MTJs to the latch's Q and/or QN outputs. This may happen automatically when power is first applied to the circuit, or it may require the assertion of one or more control signals to accomplish the task.

All input signals should use typical CMOS switching levels. Output signals should be rail-to-rail, with quiescent voltage levels essentially equal to V_{DD} and ground.

2.1.2 Power Consumption

As a general rule, power consumption should be minimized. The energy required to store a bit of information in an embedded magnetic latch should compare favorably to other non-volatile memory technologies, such as floating-gate or chalcogenide (phase change) devices. However, for applications where the MTJs are programmed only occasionally, perhaps just once before being inserted into the application, power consumption during programming may

be of secondary importance.

Another important issue is the peak current required. Higher peak currents will require wider conducting paths on the integrated circuit to reduce voltage drops, which will increase the size (and cost) of the final product. More power pins may be needed on the package, which again increases the final cost.

Except for the write operation, where the MTJs are reprogrammed, the latch should have negligible power consumption. Static power dissipation should be comparable to conventional CMOS latches, comprising only transistor leakage current.

2.1.3 Timing

The two significant timing characteristics of the embedded magnetic latch will be the time required to program new data into the MTJs and the time required to retrieve that data and convert it to valid logic levels. Since neither of these operations is expected to occur at the same frequency as other logic functions this is a loose constraint, and a reasonable goal would be less than 100 ns for either.

2.1.4 Reliability

The latch circuits must function and meet other specifications over an ambient temperature range of at least -20 °C to $+80$ °C, with a goal of operation over the full military temperature range of -55 °C to $+125$ °C. A power supply voltage tolerance of ±5% must be acceptable, with a goal of allowing a ±10% variation.

All circuits must be capable of continuous operation under any combination of environmental conditions for ten years. Although beyond the scope of this work, the MTJs must provide the necessary TMRR over the operational temperature range and for a period of ten years.

The long-term reliability of the write line metallization may be the Achilles' heel for magnetic

memories. Again, this would be of secondary importance for applications where the memory is not reprogrammed during a mission.

2.1.5 Radiation Tolerance

The MTJ elements and latch circuits must not sustain permanent failures, or degradation beyond design specifications, when exposed to the natural radiation environment of space. This includes cosmic particles with an LET of at least 140 $\mathrm{MeV \cdot cm^2/mg}$ and total dose exposure greater than 1 kGy (100 krad). The latch circuits must not exhibit single-event latchup or single-event snapback under these conditions. The threshold LET for SEU must be greater than 20 $\mathrm{MeV \cdot cm^2/mg}$ with a goal of 40 $\mathrm{MeV \cdot cm^2/mg}$.

2.2 Design Constraints

2.2.1 MTJ Behavior

Magnetic Switching

In general, the free layer of the MTJ should be reoriented as quickly and with as little energy as possible while meeting the functionality, timing, and reliability requirements. The selectivity problem described on page 24 can be eliminated by using a programming structure that does not depend on row/column crosspoint addressing, which may allow the minimum writing field to be reduced. However, the circuit design must be robust, with sufficient margin to account for the unavoidable variation in the MTJs.

Bias and Temperature Effects

Almost since the first room-temperature MTJs were demonstrated, researchers have been attempting to explain how temperature and bias voltage effect the magnetoresistance ratio. The MR is typically reduced to half its peak value when measured with a bias voltage of several hundred millivolts or more, a phenomenon that was not predicted by Julliere's simple model [105, 106]. Increasing temperature has a similar effect on the junction, where the TMR might decrease by 25% when the ambient temperature increases from 25 °C to 125 °C [107].

Early studies of the bias and temperature effects concluded that the root cause was unpolarized tunneling facilitated by defect sites in the insulating barrier [108]. This has been called "two-step" or "hopping" tunneling, and allows current to flow independent of electron spin. If it is assumed that the energetic distribution of defects in the barrier is uniform, then the density of *available* defect states as a function of energy is determined by a Fermi-Dirac function, such that the availability of states increases exponentially at higher energies. Thus, as the energy of the conduction electrons increases because of increasing temperature or increased bias voltage those carriers will be more likely to find a defect in the insulator with the same energy level, and this component of tunneling current will increase. If the two-step tunneling current becomes large with respect to the spin-polarized direct tunneling then the magnetoresistance ratio will become small.

Other researchers have attempted to unify all of the relevant temperature effects [109, 110]. For example, the degree of spin polarization in the ferroelectric layers is the fundamental phenomenon considered in Julliere's model, but this polarization is temperature dependent and decreases at higher temperatures. The electrode's magnetization at the tunnel barrier interface determines the spin polarization, so like magnetization the polarization has a $T^{\frac{3}{2}}$ dependency:

$$P(T) = P_0(1 - \alpha T^{\frac{3}{2}}) \tag{2.1}$$

where P_0 is the spin polarization at $T = 0$ K and α is a material constant. Moodera et al. studied junctions at 77 K and 295 K and observed very good correlation with this model [111]. Reported values of α for $Ni_{80}Fe_{20}$ and Co electrodes are $\alpha_{NiFe} = 3 - 5 \times 10^{-5}\ K^{-3/2}$ and $\alpha_{Co} = 1 - 6 \times 10^{-6}\ K^{-3/2}$. This difference in α explains why temperature effects are not as severe in MTJs fabricated with Co electrodes. A study of thin films of $Ni_{48}Fe_{52}$ suggests that the $T^{\frac{3}{2}}$ dependency on temperature for magnetization is typical for thicker films but as the thickness is reduced to less than about 1 nm the relationship becomes more linear, reflecting a transition from 3D to 2D behavior at a thickness of about 6 monolayers [30]. Because of this, very thin films will loose their magnetization well below the bulk Curie temperature. Put another way, the Curie temperature itself is a function of film thickness, and at some point (about one monolayer) the material becomes paramagnetic and the Curie temperature falls to zero [6].

$G_T(T)$ represents the variation in direct tunneling across the junction and is approximated as

$$G_T(T) = G_0 \frac{CT}{\sin(CT)} \tag{2.2}$$

where G_0 is the conductance at $T = 0$ K and C is a material constant. Finally, the thermal variation in spin-independent "hopping" tunneling can expressed as

$$G_S(T) = ST^{\gamma} \tag{2.3}$$

The constant S is an indication of the quality of the tunnel barrier, i.e. the density of defects. The exponent γ depends on the average number of "hops" that an electron makes as it passes through the barrier and was empirically determined to be approximately $\frac{4}{3}$. Combining these three factors gives the following expressions for the maximum and minimum conductance of

the tunnel junction:

$$G_{\max}(T) = G_T(T)[1 + P(T)^2] + G_S(T) \tag{2.4}$$

$$G_{\min}(T) = G_T(T)[1 - P(T)^2] + G_S(T) \tag{2.5}$$

Recent studies have found that typical tunnel barriers are of sufficient quality that $G_S(T)$ represents less than 10% of the total conductance [110]. Also, junctions using crystalline MgO barriers are less affected by spin independent tunneling because of a lower density of localized states in the insulator [112]. The variation of direct tunneling, $G_T(T)$, is only a few percent for typical barriers over the temperature range of interest [109]. Consequently, temperature variations in magnetoresistance ratio are determined almost entirely by the corresponding variation in spin polarization.

Furthermore, it has been shown that small changes in magnetization can cause larger changes is spin polarization at the barrier surface, and much larger changes in magnetoresistance ratio [113]. This sort of model can be extended to explain bias-dependent magnetoresistance by assuming that the spin-polarized density of states is itself bias-dependent [114]. This is consistent with measurements of junction resistance as a function of bias voltage, which show that the decrease in magnetoresistance ratio is primarily due to a decrease in the antiparallel resistance value, suggesting a loss of spin polarization [32].

While the variation in the magnetoresistance ratio is of most concern, it is important to note that the absolute resistivity of a tunnel junction is also affected by bias and temperature. Simmons proposed that the variation of the tunnel barrier's dielectric constant (relative permittivity) with temperature would have a significant impact on the tunnel resistivity, although his work did not specifically address magnetoresistance ratios [12]. Although the bulk permittivity of some insulators, such as $SrTiO_3$, can vary by two orders of magnitude from 20 K to 300 K, the permittivity of both MgO and Al_2O_3 changes by less than 1% over the same range [115]. Measurements on squares of Al_2O_3 that were 0.5 mm thick found the

temperature coefficient of permittivity to be $11.0 \times 10^{-5}/\mathrm{K}$, which also suggests that this effect is negligible for the normal operating range of electronic components [116].

Capacitance

Although it is the resistive properties of an MTJ that make it suitable for information storage, the capacitance of the junction is relatively large and cannot be neglected. The tunneling barrier is just a few nanometers thick, on the same order as the transistor gate oxide, with a relative permittivity that is at least twice that of SiO_2. The area of the junction may also be large compared to a nominal transistor, so the net result is that the MTJ accounts for one of the largest parasitic (unintended) capacitances in the circuit.

A simple parallel-plate approximation can be used to estimate the junction capacitance, but this does not include the fringing field capacitance between the edges of the free and pinned ferromagnetic layers. In practice, the width or diameter of the junction is much greater than the insulator thickness. For circular electrodes the total capacitance from the fringing field can be estimated using a "thin disk" approximation [117]:

$$C = C_{PP}\left(1 + \frac{2x_o}{\pi R}\ln\left(\frac{2e\pi R}{x_o}\right)\right) \tag{2.6}$$

where C_{PP} is the parallel-plate capacitance, R is the radius of the disk, and x_o is the insulator thickness. For a radius of 0.5 μm and oxide thickness of 3.5 nm this equation simplifies to $C = 1.035\,C_{PP}$ and the contribution due to fringing is negligible for practical purposes.

However, the thin disk approximation assumes that one of the electrodes is an ideal, infinitesimally thin disk. This is certainly not the case in reality and may significantly underestimate the fringing effect. As an example, the capacitance of a rectangular plate over a

conducting plane can be estimated as

$$C = \varepsilon_o \varepsilon_r \left(1.15 \left(\frac{W \times L}{x_o} \right) + 1.40(2W + 2L) \left(\frac{T}{x_o} \right)^{0.222} + 4.12 x_o \left(\frac{T}{x_o} \right)^{0.728} \right) \tag{2.7}$$

where ε_o is the permittivity of free space (8.854×10^{-12} F/m), ε_r is the relative permittivity of the insulator, W and L are the dimensions of the rectangle ($W \leq L$), T is the thickness of the plate, and x_o is the insulator thickness [118]. The second term in Equation 2.7 represents the contribution from the circumference of the plate and the third term is the added capacitance of the four corners. The error of this approximation is said to be less than 10% for $0.4 < T/x_o < 10$ and $0.5 < W/x_o < 40$. For $x_o = 3.5$ nm, $W = L = 40x_o = 140$ nm and $T = 10x_o = 35$ nm, equation 2.7 yields

$$C = \varepsilon_o \varepsilon_r (6.44 \times 10^{-6} + 1.31 \times 10^{-6} + 0.077 \times 10^{-6}) \tag{2.8}$$

$$= 7.827 \times 10^{-6} \varepsilon_o \varepsilon_r \tag{2.9}$$

$$= 1.4\, C_{PP} \tag{2.10}$$

which shows a significant contribution from fringing. For comparison we can use the thin disk approximation of equation 2.6 for a disk with the same area as this plate ($R = 79$ nm) to obtain $C = 1.17\, C_{PP}$, or for a disk with the same circumference ($R = 89$ nm) to obtain $C = 1.15\, C_{PP}$

Neither of these approximations is well suited to MTJs, where the thickness of the electrodes may be much greater than $10\, x_o$ and the width of the electrode is much greater than $40\, x_o$. They do, however, give an order-of-magnitude estimate of the fringing effect, and it is clear that the relative contribution of the fringing capacitance will become even smaller as the width of the electrode becomes very large with respect to the insulator thickness. For typical MTJ dimensions, it appears that the fringing capacitance will not be a significant factor in designing appropriate circuits.

A magnetocapacitance effect has been observed in MTJs that are used to sense a magnetic

field, such as in magnetic recording applications. When the applied field is switched at a high frequency the junction capacitance was found to vary with the field direction and strength, and the change in capacitance was comparable to the magnetoresistance variation ($\approx 8\%$) [119]. For embedded memory applications the magnetic fields are static during a readout operation, so any magnetocapacitance would not be relevant.

2.2.2 Reliability

Radiation Effects

Since MTJs do not rely on charge as a data storage phenomenon they are relatively insensitive to TID effects. Furthermore, the tunnel current will tend to fill trapped holes, repairing the damage caused by an ionizing event. The very thin Al_2O_3 films used in MTJs have demonstrated very little damage from TID levels of above 10 kGy [120]. The degradation due to TID is small compared to temperature and manufacturing variation so it is likely that existing MTJ memory circuits will easily cope with TID effects.

Fortunately, only low bias voltages are used with MTJs, which reduces the risk of SEGR. Tests performed with heavy ions demonstrated a small decrease in the tunneling magnetoresistance and no observable change in the magnetic layers [121].

Tunnel Insulator Reliability

Retrieving information stored in MTJs requires that conflicting constraints be met simultaneously. In order to sense the resistance of the MTJ we must apply a voltage across it and allow current to pass through. For improved signal-to-noise ratios the applied voltage (and current) should be increased, but doing so will increase power consumption and bias voltages across the MTJ should be kept well below 1 V to insure long term reliability [122–126].

However, the choice of tunnel insulator material can be an important factor. In a study

comparing time-dependent dielectric breakdown (TDDB) in Al_2O_3 and MgO barriers, it was found that the lifetime of MgO barriers was at least three orders of magnitude greater than that of Al_2O_3 barriers at any given applied voltage [127]. Extrapolations from the observed failure rate indicate that 1.4 nm Al_2O_3 barriers can achieve a 10-year lifetime only if the applied voltage is less than 0.4 V while 1.6 nm MgO barriers could tolerate a voltage of 1 V for the same period.

Electromigration

The conductors that carry the write pulse current represent a significant reliability concern. To generate the most effective magnetic field possible for a given amount of current we should concentrate that current into a thin conducting sheet directly adjacent to the free ferromagnetic layer. This goal drives the design toward very high current density levels in the write coil, which will reduce the reliability of the circuit.

For continuous, long term, operation the primary failure mechanism attributable to high current density is electromigration [128]. Conduction electrons dislodge metal atoms from the conductor's crystalline structure and carry them along on the "electron wind". Variations or defects in the conductor can cause localized areas of higher current density, causing metal atoms to be swept away from the area. This localized thinning of the conductor further increases the current density, accelerating electromigration until a void forms in the conductor and current flow ceases. On the other hand, metal atoms will tend to pile up where there are areas of lower current density. These deposited atoms form hillocks that can short to other nearby conductors, or the mechanical stress created by the growing metal structure may fracture the conductor [129].

Unfortunately, it is difficult to set universal constraints on current density because the methods used to create and pattern a metal layer greatly influence its susceptibility to electromigration. The grain structure of the metal has a large effect, as the grain boundaries provide a low

energy path for the migration of metal atoms [130]. When the grain boundaries are aligned with the current flow then electromigration is much more severe. This can lead to the surprising result that when metal lines become narrower than the average grain size their resistance to electromigration starts to increase, because at this point many of grain boundaries will be perpendicular to the electron flux. Furthermore, pure metals are not generally used for integrated circuit wiring. Silicon is typically added to aluminum to improve the reliability of aluminum/silicon contacts, and alloys with a small amount of copper added show greatly reduced electromigration.

Characterizing electromigration in copper lines shows more variability than in aluminum, due largely to differences in the fabrication processes. Typically, aluminum lines are formed by creating a thin film over the entire wafer surface and then etching away the unwanted metal, a process known as subtractive etching. The aluminum lines are then covered with a protective dielectric layer. Because of difficulties in etching copper this metal is usually patterned using a *damascene* process. In this case, the dielectric layer is formed first and trenches are etched into it where copper lines are needed. A uniform layer of copper is deposited over the dielectric, which fills the trenches as well as covering the surface of the dielectric. A chemical-mechanical polishing (CMP) step then removes the excess copper, leaving only the metal that was deposited in the trenches.

The copper film may be created by chemical vapor deposition (CVD) or by electroplating and these two processes result in significantly different grain sizes, as can the width of the trench [130]. Copper can diffuse through silicon dioxide and contaminate silicon so the trenches must be lined with a barrier material "seed" layer before the copper is applied, and the choice of barrier material will influence the current carrying ability of the copper. Similarly, the copper line is capped with another barrier material. The interfaces between the copper and the barrier layers, particularly the cap layer, provide an additional path for metal atom migration. As in the case of aluminum, the addition of low concentrations of other

metals is being investigated as a means of reducing electromigration. Despite these variations, copper interconnections typically provide at least an order of magnitude increase in reliability over aluminum lines [131, 132].

In practice, design limits on current density to prevent electromigration are typically specified by the integrated circuit manufacturer. The manufacturer may not provide a rationale for a particular design rule, so it becomes difficult to assess the degree of conservatism or the failure criteria involved. In any event, typical restrictions on average current density are approximately 1.5×10^9 A/m^2 (0.15×10^6 A/c^{2}m) for operation at 125 °C in AlCu lines. For alternating currents the design limit is based on restricting Joule heating to a few degrees, which may increase the limit on r.m.s. current density by a factor of 2 to 4.

Blech effect However, the standard guidelines for avoiding electromigration were formulated for relatively long metal lines, which would be applicable to the word line and bit line of an MRAM. In an embedded MTJ memory cell the portion of the write line that has a high current density need not be much longer than the largest dimension of the MTJ itself, and electromigration follows a much different set of rules in this regime. In short lines the force exerted on the metal atoms by the electron flux may be counteracted by the opposing force from mechanical compression and concentration gradients at the cathode end, effectively halting electromigration [133]. Blech reported that for a given current density, j, there is a critical line length, L, called the *Blech length* below which there is no atomic drift, and hence no electromigration:

$$(j \times L)_c = \text{constant} \tag{2.11}$$

where the constant is determined only by material properties and temperature. For both AlCu and damascene Cu lines resistance to electromigration starts to increase significantly for lines shorter than 100 μm, and Blech's model suggests that lines shorter than about 50 μm are essentially immune to long term electromigration failure [134–136].

More recent studies of the Blech length have found that the relationship between line length and electromigration is more complex. A study of short aluminum lines found that the drift velocity cannot be linearly extrapolated to zero, and that Blech's model was optimistic for lines near the critical length [137]. This work reported that the critical constant for aluminum at 200 °C is 670 ± 120 A/cm. Applying this value to a line that is 2 μm long with a cross-section of 1 μm by 1 μm suggests a maximum current of

$$670 \frac{\text{A}}{\text{cm}} \times (1 \times 10^{-4}\,\text{cm})^2 \times \frac{1}{2 \times 10^{-4}\,\text{cm}} = 33.5 \times 10^{-3}\,\text{A} \tag{2.12}$$

Similar investigations of the Blech length in damascene copper interconnects found the critical constant to be approximately 4000 A/cm at 300 °C for 1 μm wide lines [138]. Lines just 10 μm long were said to be "immortal", with only a slight increase in resistance after 5100 hours at 300 °C carrying a current of 1.6 MA/cm^2 (16 mA/μm^2). However, when a damascene process is used the copper grain size is strongly affected by the patterned line width. This introduces a non-linear and non-monotonic variation in jL_c, such that the value of this constant was measured to be only 1500 A/cm for 0.12 μm wide, 0.25 μm thick, lines [139].

Pulsed current Most of the research on electromigration has focused on steady-state conditions, where the current density and temperature are constant. These studies provide guidance for specifying the power distribution networks in integrated circuits, which are typically the "weakest link" for electromigration. However, studies of unidirectional pulsed current show more complicated behavior and give greater insight for embedded MRAM circuits. Studies under pulsed conditions indicate a longer Blech length than for d.c. conditions. For example, pure aluminum lines demonstrated a jL_c product of 901 A/cm for continuous current and 2314 A/cm at a duty cycle of 25% [140]. In this work the jL_c product appeared to be roughly proportional to the inverse of the duty cycle for duty cycles down to

25%, suggesting that the Blech length is determined by the average current rather than the peak current. This result further suggests that damage in these short lines is repaired by diffusion when the current is removed, which should be a function of temperature in the metal. An investigation that modeled heat conduction and diffusion provides support for this conclusion but complicates the practical application of the Blech effect [141]. These simulations indicated that the Blech length is also affected by the frequency of current pulses, in addition to their duty cycle. At low frequencies the metal has time to cool during the periods without current, which retards the diffusion process, while at high frequencies the metal maintains a higher equilibrium temperature. It was also found that as the current density increases, which increases the magnitude of the temperature excursions, the Blech length tends to be reduced toward its d.c. value.

Joule Heating While electromigration damage is primarily a long-term degradation related to the *average* current density, there is also an absolute limit to the maximum current density that can be sustained in a single short pulse. In the case of a thin metal conductor bonded to an insulating substrate the failure mechanism is the fracturing of the metallic conductor due to Joule heating and mismatched coefficients of thermal expansion between the metal and substrate. For example, common aluminum conductors on silicon dioxide dielectrics will fail with a temperature increase of 300 °C, well below the melting point [142].

For very short current pulses (less than 10 ns when the line is covered in SiO_2, less than 2 ns when the dielectric is SiN) we can assume adiabatic conditions, where no significant heat is conducted away from the metal line [143, 144]. In this case it can be shown that the temperature increase in the metal is independent of the dimensions of the line but is proportional to the square of the current density. The power dissipated by Joule heating is simply

$$P = VI = I^2R \tag{2.13}$$

and so the energy deposited in the metal is

$$E = P\Delta t = I^2 R\Delta t \tag{2.14}$$

$$= (JA)^2 \left(\frac{\rho_r L}{A}\right)\Delta t \tag{2.15}$$

$$= J^2 \rho_r LA\Delta t \tag{2.16}$$

where I is the current flowing in the metal line segment, J is the current density, R is the resistance of the line segment, Δt is the duration of the current pulse, ρ_r is the resistivity of the metal, L is the length of the line segment and A is its cross section area. The temperature increase in the line segment due to this energy is

$$\Theta = \frac{E}{c_m m_m} \tag{2.17}$$

$$= \frac{E}{c_m(\rho_m LA)} \tag{2.18}$$

where Θ is the temperature increase in the line segment, c_m is the specific heat capacity of the metal, m_m is the mass of metal line segment, and ρ_m is its mass density. Combining these equations gives

$$\Theta = \frac{J^2 \rho_r LA\Delta t}{c_m \rho_m LA} \tag{2.19}$$

$$= J^2 \Delta t \left(\frac{\rho_r}{c_m \rho_m}\right) \tag{2.20}$$

The last factor in Equation 2.20 is a constant for any given metal (neglecting the change in resistivity with temperature). If ρ_r is given in $\Omega\cdot\mathrm{m}$, c_m in $\mathrm{J/kg\cdot K}$, and ρ_m in $\mathrm{kg/m^3}$ then this term has units of $\Omega\cdot\mathrm{K\cdot m^4/J}$. But an ohm is a $\mathrm{V/A}$ and a joule is an $\mathrm{A\cdot V\cdot s}$ so an Ω/J is equal to $1/\mathrm{A^2\cdot s}$ and our new material constant has the appropriate units of $\mathrm{K\cdot m^4/A^2\cdot s}$.

For aluminum, copper and gold the relevant material properties are given in Table 2.1. As an example, suppose a copper film 20 nm thick and 2 μm wide carries a current of 20 mA for

10 ns. The current density is 5×10^{11} A/m^2, and the resulting temperature increase is 12.3 K. If the same line were composed instead of aluminum or gold, the temperature increase would be 27.3 K or 24.5 K, respectively.

Property	Units	Aluminum	Copper	Gold	Tungsten
ρ_r	$\Omega \cdot$m	26.5×10^{-9}	17×10^{-9}	24.4×10^{-9}	52.8×10^{-9}
c_m	J/kg$\cdot$K	900	385	129	132
ρ_m	kg/m^3	2700	8960	19300	19250
$\rho_r/c_m\rho_m$	K$\cdot$m^4/A$^2\cdot$s	10.9×10^{-15}	4.93×10^{-15}	9.80×10^{-15}	20.8×10^{-15}

Table 2.1: Material properties for common conductors

Pulses with very high current density, short duration, and low duty cycle are not common in integrated circuits and have not been studied extensively. They most often occur in the context of electrostatic discharge (ESD) events, FPGA programming, or radiation-induced latchup. An important question related to such pulses is how the reliability of the conductor is affected by them. A study of the reliability of aluminum lines that had been subjected to a single pulse of 20 MA/cm^2 (2×10^{11} A/m^2) found that there was no long term degradation in their mean time to failure [145]. The conclusion of this study was that pulses below the catastrophic failure limits did not permanently damage the conductor.

2.3 Circuit Design

2.3.1 MTJ Simulation Models

The magnetic tunneling junction has a number of characteristics that are significant to the circuit simulation. To first order, the MTJ can be modeled as a bias-dependent resistance in parallel with a fixed capacitance. The resistance of the junction is greater when the surrounding ferromagnetic layers have antiparallel fields, but circuit simulators typically cannot model this effect so it becomes necessary to develop slightly different models for the parallel and antiparallel cases. The fixed capacitance is simply that of the two conducting

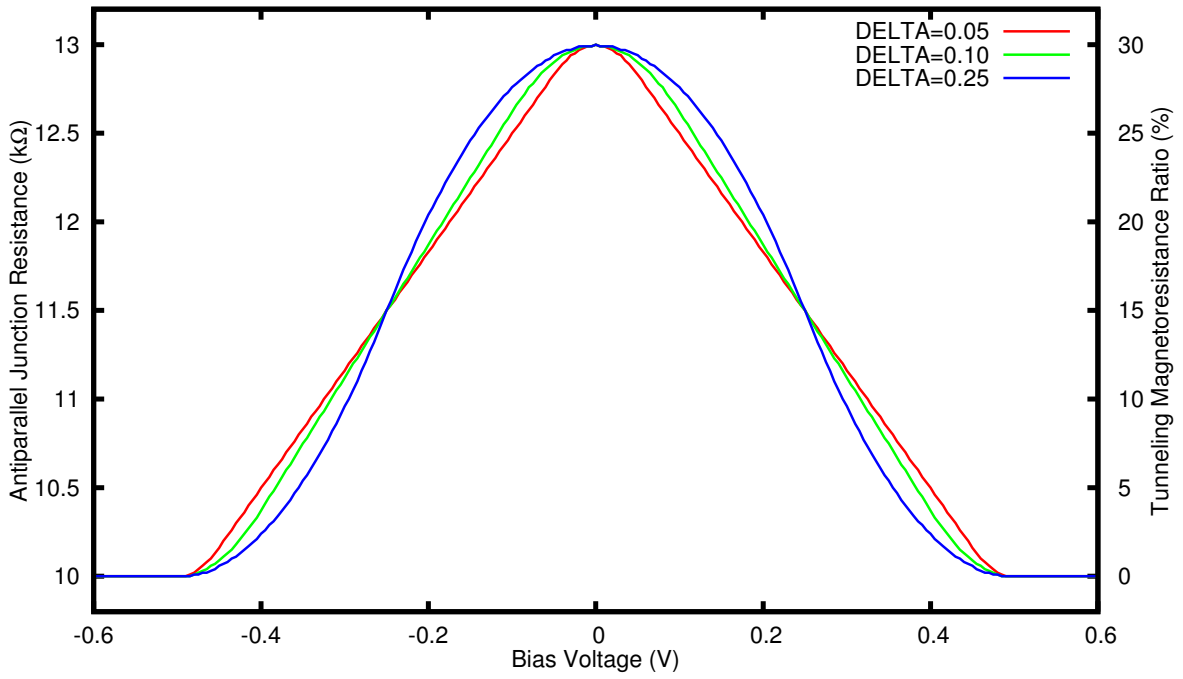

Figure 2.1: SPICE model of MRR bias dependence

ferromagnetic layers separated by the thin tunneling oxide.

The magnetoresistance ratio (MRR) is known to be a strong function of the voltage across the tunneling junctions, falling off rapidly over a range of several hundred millivolts [146–148]. For circuit simulations the resistance of the MTJ with antiparallel magnetic fields is modeled as

$$R_{AP} = \begin{cases} R_P & \text{if } V_{BIAS} > V_{MAX}, \\ R_P\left(1 + MRR(\frac{V_{MAX} - V_{BIAS}}{V_{MAX}})\right) & \text{if } V_{BIAS} < V_{MAX}. \end{cases} \tag{2.21}$$

where R_{AP} is the effective resistance of the junction, R_P is the resistance of the junction with parallel magnetic fields (i.e. the minimum resistance), V_{BIAS} is the voltage across the junction and V_{MAX} is the voltage across the junction that reduces the MRR to zero.

Figure 2.1 illustrates how the bias dependence is modeled in HSPICE, a popular transistor-level simulator used in integrated circuit design. The MTJ with antiparallel fields is represented as a voltage-controlled resistor with a piecewise-linear transfer function. In this example R_P is 10 kΩ, the MRR is 30% and V_{MAX} is 500 mV. The model also has a smoothing

parameter, DELTA, which forces the transfer function to be smoothly continuous in order to prevent instability and lack of convergence in the simulations. See Appendix B for a listing of the HSPICE model file.

Assuming that the ferromagnetic layers represent ideal conductors, the parallel-plate capacitance of the junction can be given as the oxide capacitance

$$C_o = \frac{\varepsilon_r \varepsilon_o}{x_o} \tag{2.22}$$

where ε_o is the permittivity of free space, ε_r is the relative permittivity of the insulator, x_o is the thickness of the oxide and C_o is the capacitance of the oxide per unit area [149]. Reported values of ε_r for common MTJ insulators include 8.0 [22] and 9.4 [115] for Al_2O_3 and 9.8 for MgO [115]. For a typical oxide thickness of 3.5 nm the parallel-plate capacitance is therefore $\approx$20 fF/μm^2. As discussed in Section 2.2.1, the fringing capacitance of the junction is a small but non-negligible fraction of the total. However, the area, shape, and thickness of the tunneling barrier to be used with the circuits discussed here are not yet determined and these factors contribute a much larger uncertainty than an imprecise estimate of the fringing capacitance. Consequently, the characteristics of the circuits will be evaluated over a range of junction capacitance values rather than attempting to calculate a precise, but completely wrong, value for the fringing effect.

Previously developed SPICE models for MTJs were primarily concerned with modeling the bias-dependent magnetoresistance, either with a VCR [150] or by using the transconductance of a MOSFET [151]. The earlier model by Das and Black included the ability to switch the state of the MTJ from parallel to antiparallel by applying the appropriate current to the "write line" terminals, but this would seem to be an unnecessary complication for most purposes. Neither of these models included the junction capacitance of the MTJ itself, perhaps because they were intended for MRAM simulations where the MTJ capacitance was dwarfed by the capacitance of the write line and bit line.

2.3.2 Prior MTJ Latch Designs

Typical magnetoresistance ratios are smaller than the normal spread in circuit parameters due to manufacturing variations and operating conditions so it is not possible to reliably sense the MTJ's state (parallel or antiparallel) by measuring its absolute resistance. Therefore, practical approaches to magnetic memories commonly use a reference circuit that represents a resistance equal to a midpoint between the two possible MTJ states [152, 153]. An alternative scheme employs two MTJs with opposite states in each memory bit and compares them [154–156].

The use of two counter-programmed MTJs in a simple latch was suggested by Daughton in 1997 [22], and later patented [157], with the simple circuits shown in Fig. 2.2. These circuits

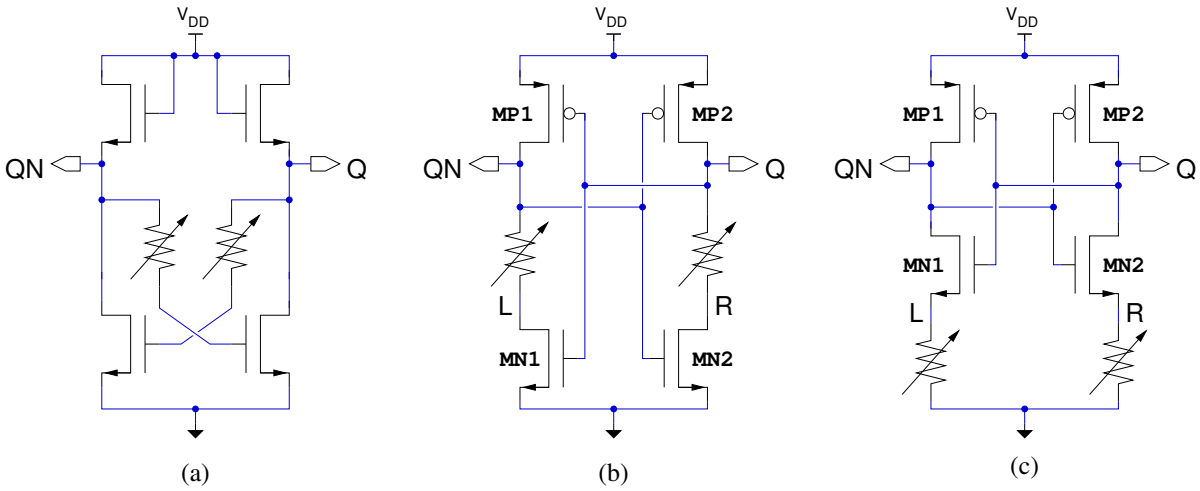

Figure 2.2: Daughton's latches [157]

were designed to assume a logic state that was predetermined by the two MTJs, one parallel and the other antiparallel, when power was first applied. The circuit shown in Fig. 2.2(a) is not suitable for embedded latches for two reasons. First, using NMOS transistors as high side loads (pull-ups) means that the circuit outputs cannot rise to V_{DD} and the maximum voltage for a logic '1' output will be V_{DD}-V_{TN}. Second, the circuit has significant d.c. power consumption because the NMOS loads are always conducting.

The second Daughton latch, in Fig. 2.2(b), does not have these disadvantages and may be suitable for commercial applications. However, it does not offer significant immunity to single event upsets, as shown in the simulation results in Fig. 2.3. These simulations used transistor models from a commercial 0.18-µm bulk CMOS process operating at 1.8 V. The

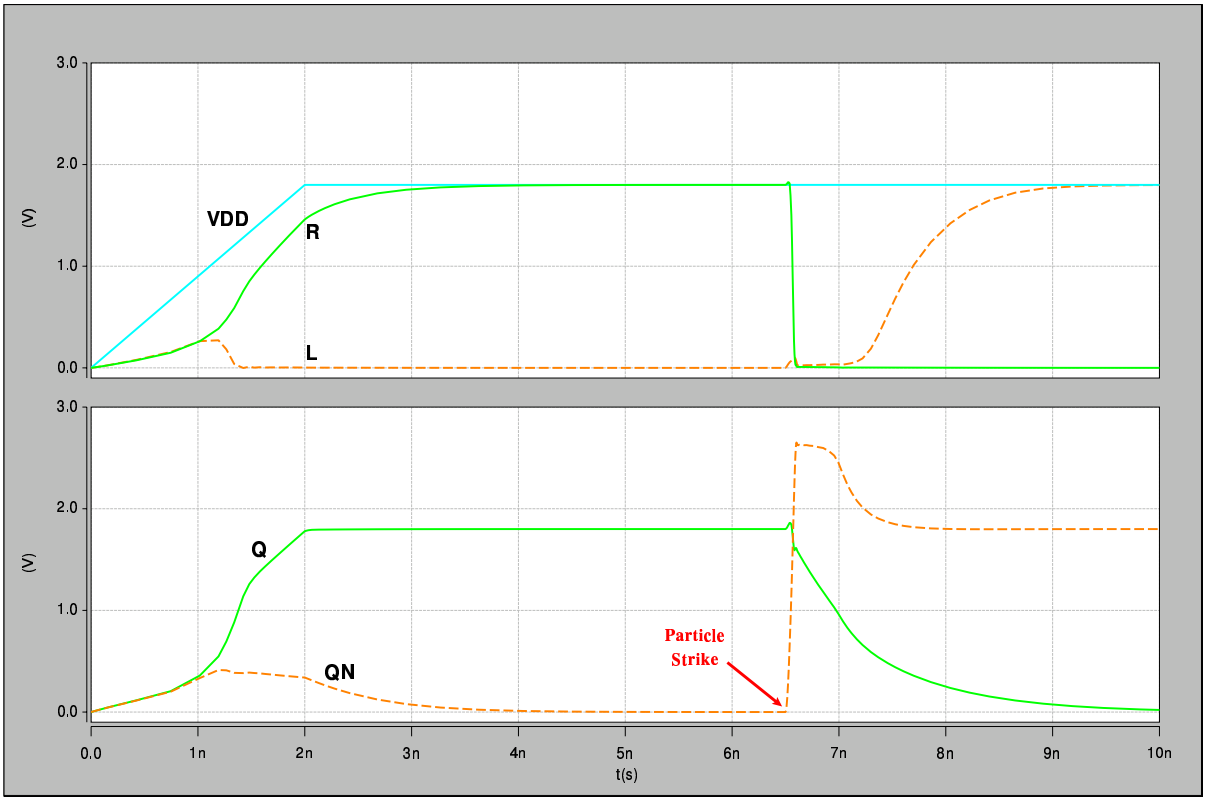

Figure 2.3: Daughton latch (b) simulation with particle strike

W/L ratio (in µm) for the PMOS transistors is 5.0/0.18 and for the NMOS transistors it is 3.0/0.18. The resistance of the parallel MTJ is 30 kΩ and the TMRR is 30%. As shown in Fig. 2.3 the latch does indeed resolve the state of the MTJs and enter the desired logic state as V_{DD} rises. A current pulse with a total charge of 1 pC is injected at node QN at 6.5 ns to simulate a cosmic particle strike. QN immediately rises until the drain-to-body diode in `MP1` conducts and clamps the voltage at about 2.6 V. Since QN is connected to the gates of `MP2` and `MN2`, `MP2` ceases to conduct current while `MN2` is enabled, causing node R to fall very quickly while Q is discharged more slowly through an MTJ. At this point, just after the strike, two of the four nodes in the latch have been forced to incorrect logic levels. Both L and R are

close to ground so any information stored at these nodes is lost, and the correct state of the latch can not be discerned using them. Examining nodes QN and Q shows that QN is actually at a higher voltage than Q, so the latch resolves to a stable but incorrect state where QN is high and Q is low. The MRR of the MTJs has little influence on this behavior. Increasing the nominal resistance in an attempt to increase the isolation between the nodes is not effective, and only slows the transition of Q to the upset state.

The latch shown in Fig. 2.2(c) has been refined and developed by several researchers, including the version suggested by Wang et al. [158] shown on the left in Fig. 2.4. Note that

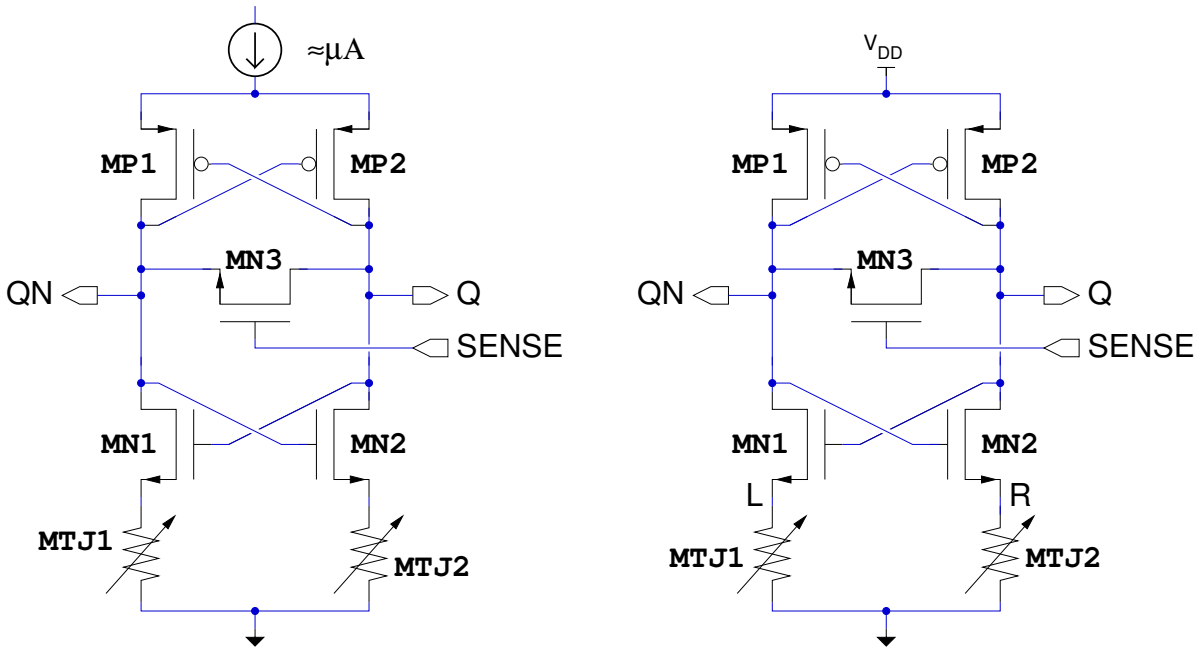

Figure 2.4: Wang Latch [158] and Unbalanced Flip Flop [159–161]

an additional transistor, MN3, has been added to enable the latch to be reset at any time. This transistor equilibrates the cell to a metastable state when SENSE is asserted, allowing the MTJs to determine the quiescent state when SENSE is then deasserted. The published examples of this circuit typically use an NMOS transistor for MN3, which would be effective only if the metastable voltage on Q and QN was relatively low, implying that the resistance of the MTJs is small compared to the transistor impedance. It seems likely that in most practical cases a PMOS transistor would be better suited to this function.

For the latch shown on the left in Fig. 2.4 the need for a constant current source complicates the circuit somewhat and results in logic '1' levels that do not reach the positive supply voltage. A slightly simplified version of this latch that also has rail-to-rail outputs has been proposed by several authors for use in programmable logic structures and is shown on the right in Fig. 2.4 [159–161]. One author [161] calls this circuit an "unbalanced flip flop" (UFF), because the MTJs have the purpose of intentionally unbalancing the cell and introducing a preferred state.

Although the authors of one paper point out that the "MTJ state is not sensitive to SEU" [161], this does not imply that the state of the flip flop itself is also insensitive to cosmic particles. In fact, simulations show that in this configuration the MTJs provide no significant protection from single event upsets, as shown in Fig. 2.5. The simulation conditions here are

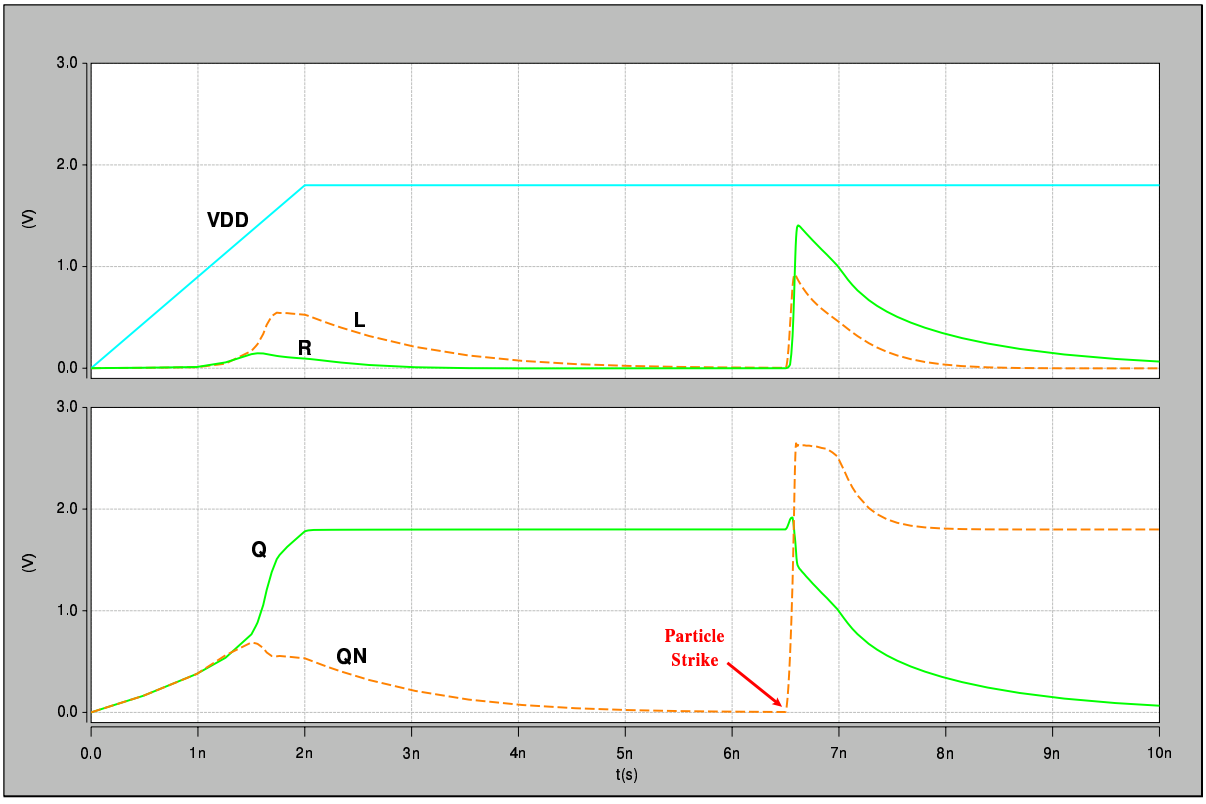

Figure 2.5: UFF simulation with particle strike

the same as those for Fig. 2.3 above. The SENSE signal and MN3 are not used to reset the latch

in this simulation, instead the latch resolves the MTJ states as power is applied. Note that when the latch is in its quiescent state both R and L are at ground so there is no information stored in the state of these nodes. When current is injected into QN this node rises quickly. Node L is pulled up through MN1 and R also rises because of capacitive coupling through the gate of MN2. Since MN2 is now enabled and MP2 is disabled, node Q begins to discharge through MN2 and MTJ2. As with Daughton's other latch we see that QN is now at a significantly higher voltage than Q and the latch resolves to an incorrect state.

It should also be noted that during the SENSE phase, when the latches of Fig. 2.4 are being precharged and equilibrated, that the voltage across the MTJs is equal to the supply voltage less the transistor threshold voltages, or V_{DD}-V_{TP}-V_{TN}. There is a risk of permanent damage to the MTJs if this voltage exceeds ≈1 V (see Section 2.2.2).

Chapter 3

Differential MTJ Latch

In the course of this research two proposed latch designs that combine MTJs with SEE-resistant circuits for embedded magnetic memory cells have been developed. The first is a differential sensing latch that uses the MTJ resistance to provide SEE resistance [162–165]. High-value resistors have been used to increase SEE immunity in a variety of circuits [88, 166] but this work is novel in that the resistors are also used for non-volatile storage.

The circuitry needed to read the cell's logic state is shown in Fig. 3.1. The two MTJ elements are shown as variable resistances `MTJ1` and `MTJ2`. Above the MTJs are two cross-coupled

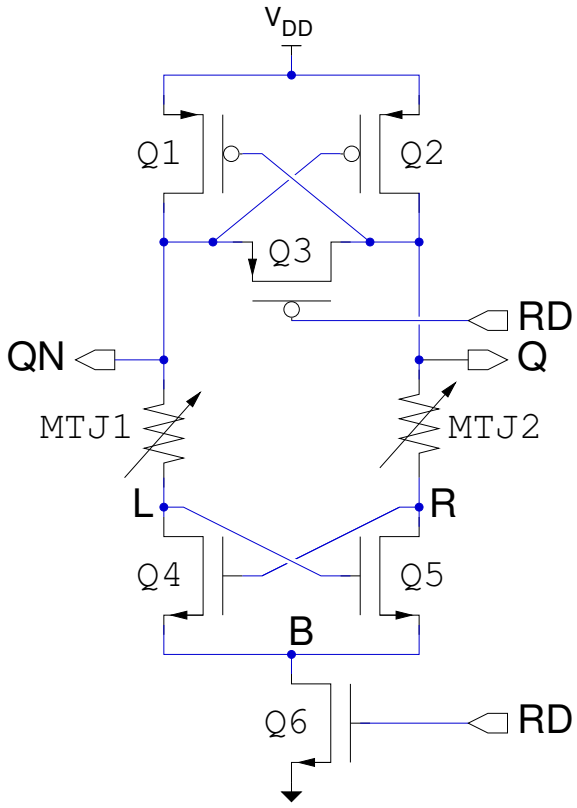

Figure 3.1: Dual-MTJ latch cell schematic

PMOS transistors, Q1 and Q2, and below them are two cross-coupled NMOS transistors, Q4 and Q5. These four transistors comprise a static latch. If the MTJs were replaced by wires the four transistors would then form the simple two inverter latch that is the heart of CMOS static RAM cells. However, in this cell the desired state is written magnetically, rather than electrically, by altering the magnetoresistance of the MTJs. Transistors Q3 and Q6, along with the signal RD, allow the difference in resistance between MTJ1 and MTJ2 to set the state of the static latch formed by the cross coupled transistors. The true and complemented logic levels corresponding to the stored state are then available on the Q and QN output signals.

The cell's read operation has been simulated in HSPICE, as shown in Fig. 3.2. Transistor parameters were for a commercial 0.18-μm CMOS process. All PMOS transistors have a W/L ratio (in μm) of 5.0/0.18, Q4 and Q5 are 3.0/0.18 and Q6 is 0.5/0.5. The TMR is assumed to be 30% with a parallel resistance of 20 kΩ. The nominal supply voltage is 1.8 V.

Initially, Q is low and QN is high, so the cross-coupled transistors are storing a logic '0' value. RD is initially asserted so Q3 is disabled while Q6 is conducting. Very little current flows through the MTJs under these conditions, so nodes L and QN are at the same voltage and nodes R and Q are at the same voltage. When RD is deasserted at 1 ns Q3 becomes conducting while Q6 is disabled and all nodes begin to rise. Q1 and Q2 will source current until Q and QN rise to V_{DD}-V_{TP}. , where V_{DD} is the supply voltage and V_{TP} is the threshold voltage of the PMOS transistors. Similarly, Q4 and Q5 will allow node B to rise to approximately V_{DD}-V_{TP}-V_{TN}. Q3 effectively shorts Q and QN together so they will rise together. Without Q3 a much longer time period would be required to equalize these voltages, which is essential to resolving the small difference in MTJ resistance.

When RD is asserted again at about 4 ns Q3 is disabled while Q6 is enabled, bringing node B to ground. Now the gate voltage, V_{GS}, of Q4 and Q5 is much greater than their threshold voltage so these transistors start conducting current, limited by the MTJ resistances. In this simulation MTJ1 has a the higher, antiparallel, resistance so less current flows into node L and this node

falls slightly faster than node R. As L falls it tends to reduce the V_{GS} of Q5 so this transistor will sink less current, reinforcing the ΔR of the MTJs with positive feedback to keep node R high. Since Q4 remains active it will begin to sink current through MTJ1 and start to pull node QN low, whereas little current flows through MTJ2 because of the low V_{GS} at Q5. As QN falls, Q2 conducts more current and brings node Q up to V_{DD}. As node Q rises Q1 will be disabled so the cross-coupled PMOS transistors also provide positive feedback to reinforce ΔR. All the nodes in the cell soon rise to V_{DD} or fall to ground, and the cell achieves a stable state with virtually no static power dissipation and no voltage across the MTJs. Note that in a quiescent state, during either a read or write operation, there is no voltage across the MTJs. While switching from one state to the other a voltage is applied to the MTJs but this is typically less than 1 V and has a duration of less than a nanosecond.

The critical single-event hardening of the proposed memory cell is achieved by serendipitous use of the high resistance of the MTJs themselves to provide decoupling, so that nodes Q, QN, L and R become four independent and redundant storage nodes. The simulation in Fig. 3.2 illustrates how the memory cell recovers from a cosmic particle strike to a PMOS transistor. In this simulation the drain of Q1, connected to node QN, receives an injected charge of 1 pC at 6.5 ns. QN quickly rises from near ground to above V_{DD}. Since QN is also connected to the gate of Q2 the disturbance is capacitively couple to Q, but this coupling does not move Q toward an invalid state. The high voltage on QN also reduces V_{GS} on Q2 so that this transistor is no longer able to source current into Q. Fortunately, the correct level is maintained on Q because of the parasitic capacitance at this node. The resistance of MTJ1 isolates node L from the single event at QN, which is critical because if the voltage on L rises sufficiently that the V_{GS} of transistor Q5 becomes greater than V_{TN} then this transistor will sink current (unopposed by Q2) and pull nodes R and Q low. If nodes R and Q should be pulled low before QN returns to its correct state then an SEU will occur. However, in this simulation node L remains low and nodes R and Q remain high so the feedback from these three valid nodes is able to restore the correct state on node QN. Within about 2 ns the deposited charge is

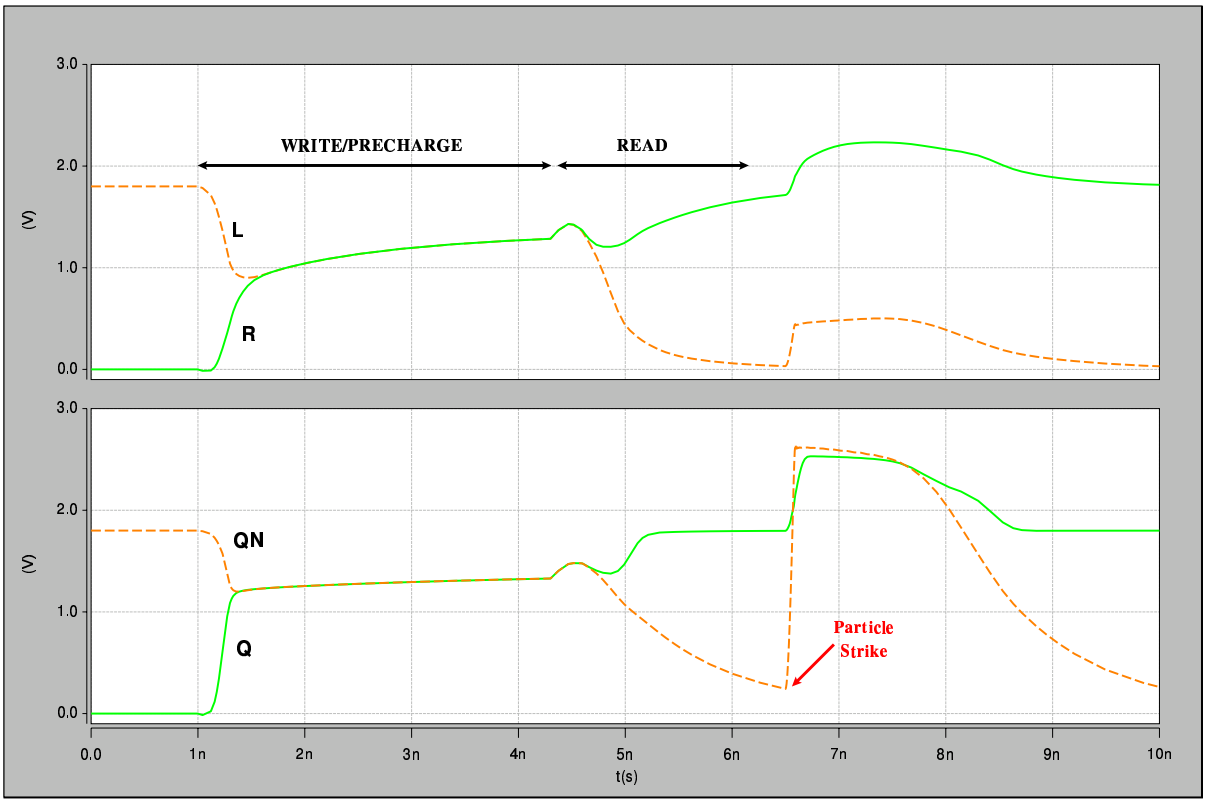

Figure 3.2: Simulated cosmic particle strike to PMOS of dual-MTJ cell

dissipated and the cell has fully recovered.

Although these simulations were performed a TMRR of 30%, the latch circuit can work with lower values if the parasitic capacitances are carefully balanced and a longer read time is allowed. Higher values of resistance yield improved immunity to cosmic particles at the expense of longer read times.

The resistance value of the MTJs is related to the radiation tolerance of the cell in two ways. A high resistance value increases the decoupling between the NMOS and PMOS transistors, so that the cell is immune to larger voltage transients from heavier or more energetic cosmic particles. Unfortunately, the correct state of the cell is restored by current flowing through a MTJ so larger resistance values also will delay the recovery of the cell. The parallel capacitance of the MTJ is also a factor since the single event transients tend to have very fast edges, and it is important to carefully model the MTJ and all parasitic capacitances in the cell.

The differential latch cell shares one characteristic with other prior art designs that may limit its usefulness in a cosmic ray environment. The QN and L signals from Fig. 3.2 are reproduced in Fig. 3.3, along with a plot of their difference. Note that the difference between

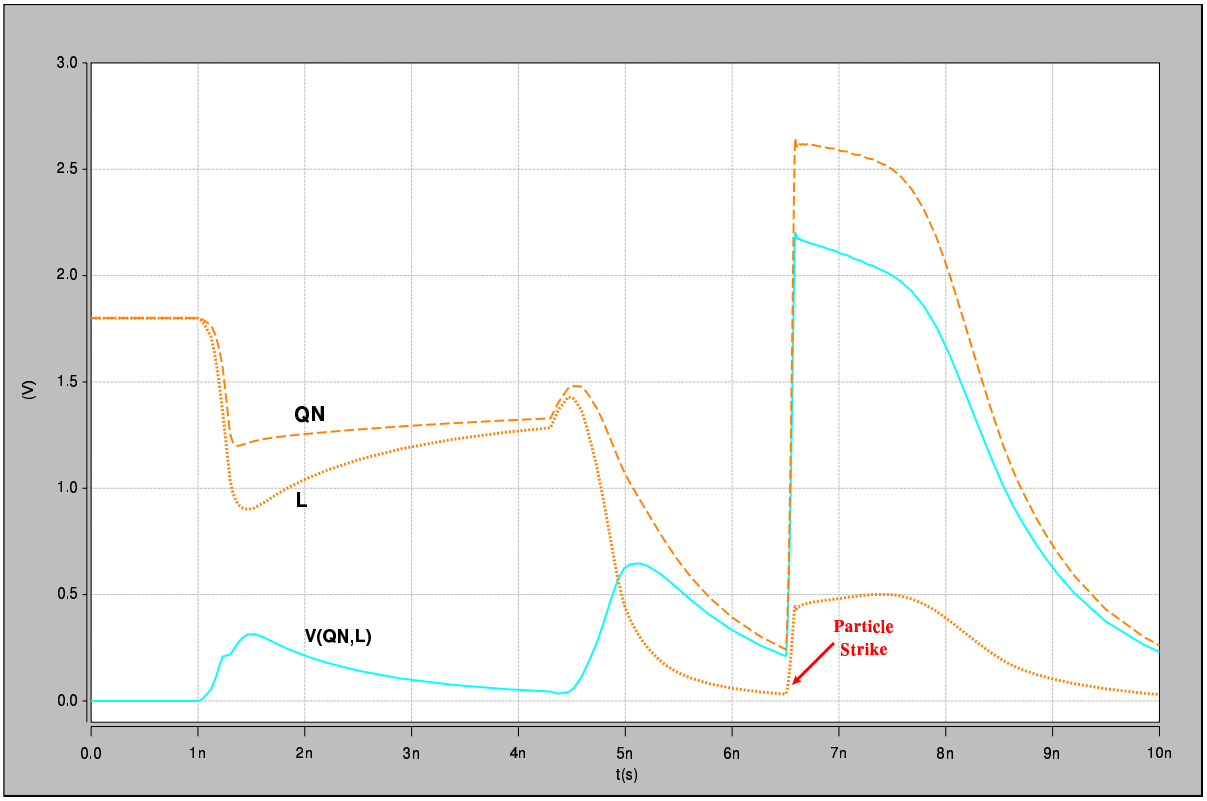

Figure 3.3: MTJ bias during cosmic particle strike

these nodes is just the bias across `MTJ1`. During the latch's recovery from the particle strike there is a period of about 1 ns when the MTJ bias exceeds 1 V. We know from studies of dielectric breakdown that long term exposure to voltages in this range can permanently damage thin insulating films, but the effect of extremely short pulses is unknown. Unfortunately, this risk is present in any circuit where an MTJ is connected to a transistor source/drain diffusion that may be reverse biased during operation, which includes all of the designs described in Section 2.3.2.

A second undesirable characteristic of the differential latch is that there is no common connection for the two MTJs. As will be discussed in Chapter 5, it is highly desirable to have

both MTJs in physical and electrical contact with the conductor that serves as the write line. Without this common node it becomes necessary to add one or more transistors to convey the write current between the MTJs. Although Q3 could conceivably fill this role, the very large size of the required pass transistor is discouraging.

In summary, the differential latch successfully extends prior art into circuits that have a high degree of immunity to single event upset. Although it shows promise in meeting the goals for this research, the disadvantages described above motivated the search for a superior circuit. For this reason the evaluation of the differential latch was rather limited in scope.

Chapter 4

MTJ Shadow Latch

The second proposed latch design is a magnetic shadow latch and takes advantage of an independent double-gate (IDG) CMOS technology developed at American Semiconductor, Inc. A shadow latch integrates a non-volatile storage cell into a conventional latch cell, so that during normal operation the memory cell has the superior speed and power characteristics of a static latch. When desired, the contents of the static latch can be transferred into the non-volatile shadow memory for retrieval at a later time. This capability is particularly attractive for circumvention applications, so that the operating state of the device can be easily restored after a nuclear event.

4.1 Double-gate Transistors

American Semiconductor's IDG transistors are called *FlexFETs* because the double-gate structure allows a variety of design goals to be pursued by changing how the bottom gate is connected. Combining MTJs with FlexFETs indeed provides unique circuit design opportunities [167–169].

The FlexFET manufacturing process begins with wafers that have a thin ($\approx$ 140 nm) bottom oxide layer, or BOX, underneath a thin ($\approx$ 180 nm) layer of high-resistivity silicon [168–170]. A 120 nm layer of heavily doped silicon is added to serve as the source and drain material. The transistor channel region is etched as a trench through the source/drain silicon, and this layer is also removed where needed to isolate transistors from each other. A heavy, retrograde implant then forms the self-aligned bottom gate under the channel. A thin gate oxide follows,

and then the top gate layer of TiN/polysilicon is added and patterned. An STI (shallow trench isolation) dielectric is deposited in the areas between transistors and the surface is planarized. Trenches in the STI and source/drain regions are then etched, and the bottom layer of metal interconnect is added by a damascene process, with the result shown notionally in Fig. 4.1. This completes the transistor level processing, and additional layers of metal interconnect are then added using common techniques and materials.

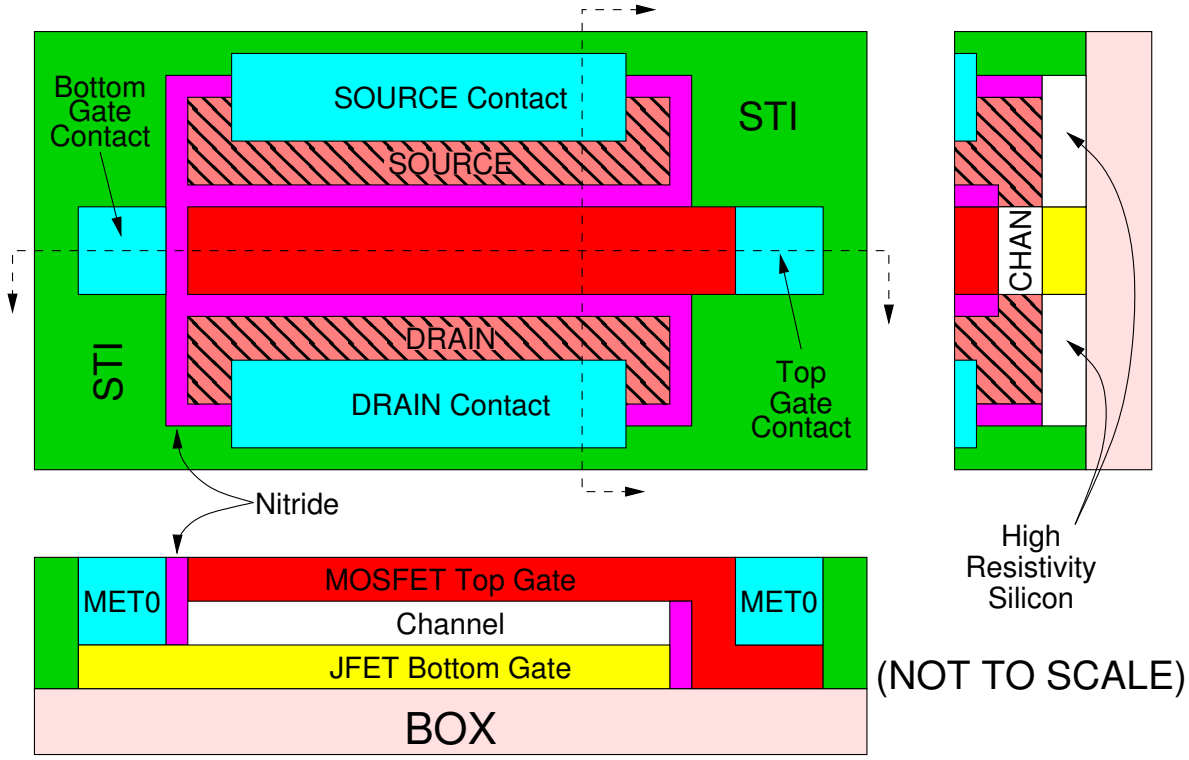

Figure 4.1: FlexFET cross section [168]

The FlexFET transistors are true four-terminal devices. In conventional MOS transistors the fourth terminal would be the transistor body or substrate, but in this case the FlexFET functions like a parallel MOSFET and JFET where the fourth terminal is the JFET (bottom) gate. The bottom gate can be connected to the transistor source to allow the FlexFET to be used as a three-terminal MOSFET, but preserving the distinct bottom gate connection introduces an additional degree of freedom for the circuit designer [171]. Altering the bottom gate voltage causes a nearly proportional change in the MOSFET threshold voltage, which in

turn varies the drain current. For example, changing the bottom gate voltage from 0.0 V to 0.5 V decreases the MOSFET (top gate) threshold from 0.425 V to 0.0 V, a much larger shift than could be accomplished using the MOSFET body effect [172]. Of course, since the bottom gate is a PN junction the circuit designer must avoid forward-biasing. The FlexFET technology has also demonstrated a high level of radiation tolerance, which complements the inherent tolerance of the MTJ [168, 173].

4.2 Shadow Latch Circuit

A simplified schematic of the information-storing circuitry of the shadow latch cell is shown in Fig. 4.2, where bottom gates are connected to the transistor source unless explicitly drawn otherwise [165, 174]. Because the achievable magnetoresistance ratio is smaller than the

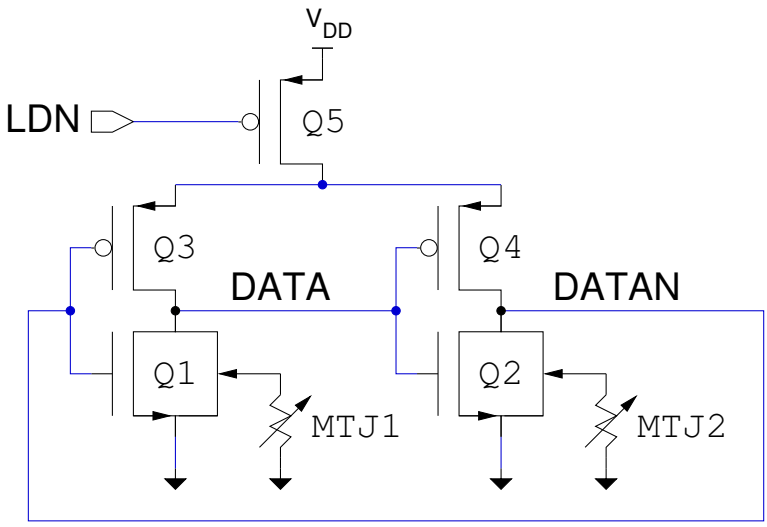

Figure 4.2: Shadow RAM cell

normal variation in circuit elements due to manufacturing and the environment, two MTJs programmed to opposite magnetic states (one parallel and the other antiparallel) are used. Differential sensing can then determine the relative resistance of the two MTJs. The MTJs are shown as variable resistances MTJ1 and MTJ2 and are connected to the bottom gate of NMOS

transistors Q1 and Q2. With PMOS transistors Q3 and Q4 the circuit strongly resembles a conventional cross-coupled inverter RAM cell.

In the interest of clarifying the role of the MTJ in the latch, assume for the moment that the internal nodes of the cell have been discharged to ground. When the load signal, LDN, is high Q5 is turned off and all circuit nodes remain discharged (see Fig. 4.3). When LDN falls Q5 is

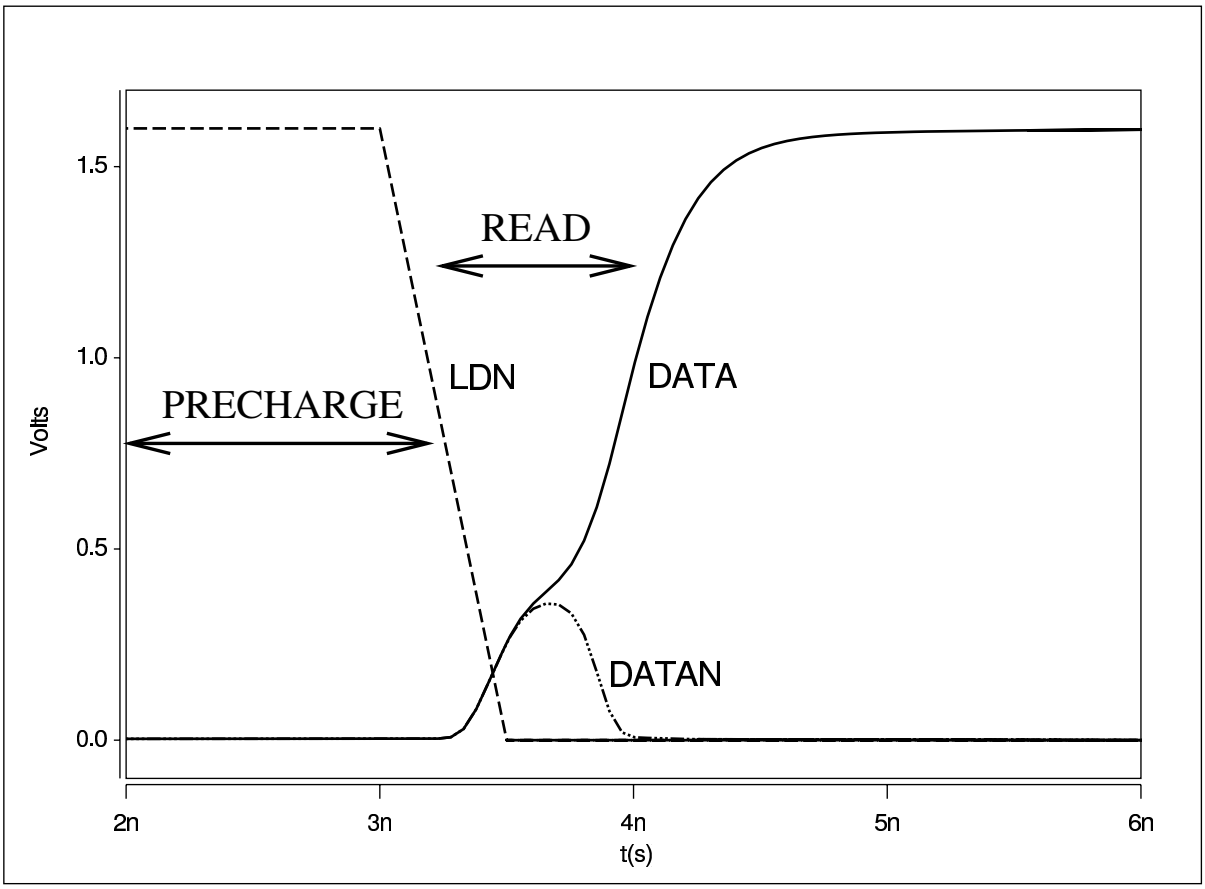

Figure 4.3: SPICE simulations

enabled, pulling the sources of Q3 and Q4 high and enabling them as well. The bottom gates of Q1 and Q2 are also pulled high by capacitive coupling to the top gate and drain, then begin to fall as their capacitance is discharged through the MTJs. The MTJ with lower resistance will pull its associated transistor's bottom gate down faster, giving this transistor less drive than the other NMOS and unbalancing the cell. Regenerative feedback causes the cell to snap to the desired state, where there is no d.c. power consumption and no voltage across the

delicate tunnel insulators.

The same approach can be applied to a single-event resistant latch design, such as the DICE cell, if greater immunity to cosmic rays is needed [95]. The schematic and layout for a magnetic shadow DICE cell is shown in Fig. 4.4, where the placement of transistors and interconnects in the schematic mimics their corresponding locations in the layout. (The SPICE netlist is provided in Appendix D.) The write line has been omitted to make the transistor layout more clear, but it is shown later in Fig. 5.3. The local interconnect in Fig. 4.4 uses the bottom two metal layers, METAL0 and METAL1, and the write line passes over the latch in the next layer above, METAL2. The patterning of the MTJs is not part of this layout but their intended location is indicated. Note that there are no edges or other features in the layers below the MTJs that might reduce the planarity or smoothness of the foundation for the magnetic layers.

The essential core of the DICE cell is formed from four complementary pairs of transistors, as shown earlier in Fig. 1.28. Transistors Q5P through Q8P and Q5N through Q8N form this core in Fig. 4.4, with the functional latch connections made using the top gates of each FlexFET. The MTJs are connected to the bottom gates of the NMOS transistors, with MTJ1 connected to Q6N and Q8N at node BG13 while MTJ2 is connected to Q5N and Q7N at node BG02. The terminals of the MTJs that are disconnected in this schematic correspond to the top surface of each element, which is connected to the Metal2 write line that passes over the top of it (see Fig. 5.3).

To make a usable DICE latch it is necessary to add output drivers to provide higher current drive levels and to prevent charge sharing with the internal storage nodes. Since the cell will be required to sense relatively small differences in the resistance of the MTJs it is critical to provide the same amount of parasitic capacitance loading at each internal node, so four identical output drivers are used. Note that connecting internal latch nodes together, so that only two output drivers are needed, would destroy the SEU immunity of the cell. The drivers

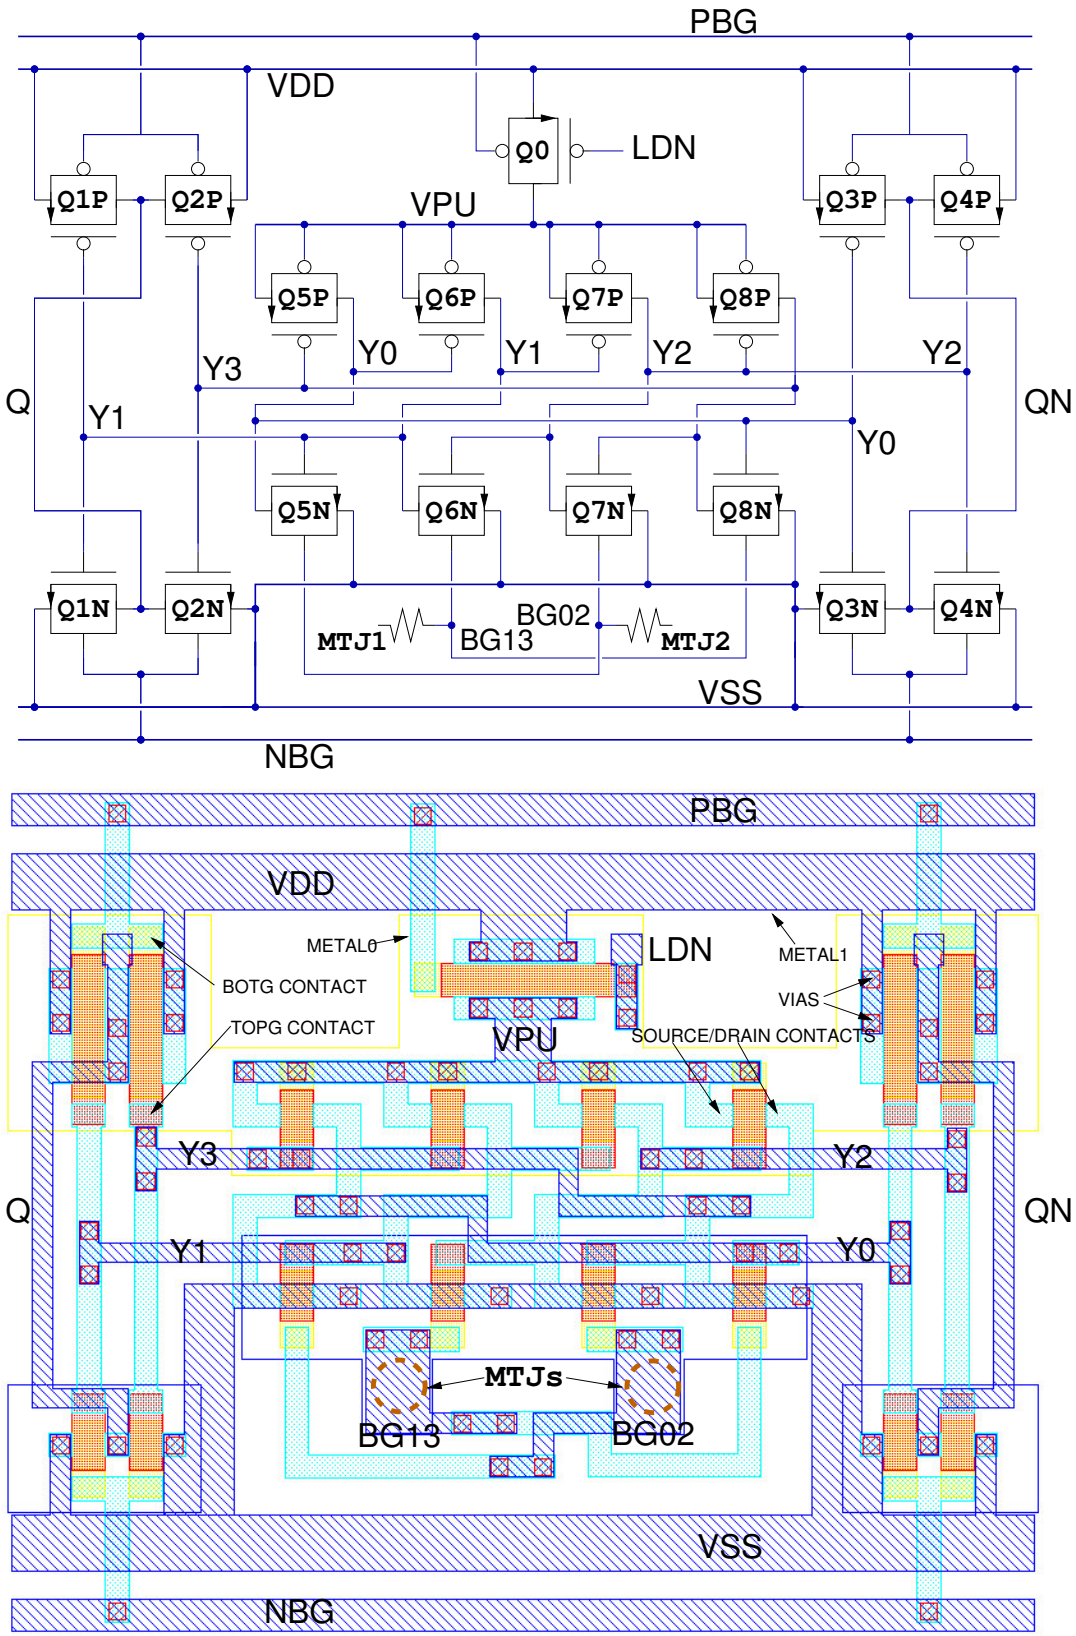

Figure 4.4: Shadow RAM cell layout

shown in Fig. 4.4 are simple inverters formed with PMOS transistors Q1P through Q4P and NMOS transistors Q1N through Q4N. The outputs of each pair of inverters are tied together to create the Q and QN outputs from the latch. The designation of one output as the true data and the other as the inverted data is arbitrary, but must be consistent with the definition of the polarity of the input data that programs the MTJs.

Two important design features of the latch have been intentionally neglected in the discussion above, because they are influenced by the nature of the embedded application and are somewhat independent of the central theme of this research. The first of these is the data input circuitry for the DICE latch, which allows a logic value to the stored in the latch independent of the stored state in the MTJs. This is the conventional method of loading the DICE latch, and the inventors suggested the two methods shown in Fig. 4.5 [95]. The circuit shown at the top of the figure uses NMOS pass transistors to load data, which is the simplest method but requires careful design to avoid charge sharing and accommodate the data-dependent loading seen at the DATA input. With additional transistors these disadvantages can be removed, as shown in the bottom circuit. Using clocked inverters instead of pass transistors eliminates any charge sharing and presents a fixed capacitive loading at the circuit inputs. A common variation on the pass transistor approach is to use four pass transistors, two driven by DATA and two driven by its inverse [96, 175]. Using four pass transistors also tends to equalize the loading on the four internal DICE nodes, which is a decided advantage for the magnetic shadow RAM.

The second design feature is the precharging (or more correctly, predischarging) of the latch. The success of the circuit in resolving small differences in MTJ resistance depends upon the equilibration of the internal nodes. There are three ways in which the magnetic shadow latch might be used in an application, and the requirements for precharging the cell are unique to each of these. The first case is if the latch is required to determine the state of the MTJs only when power is first applied, and then all nodes are initially at ground and no overt action is

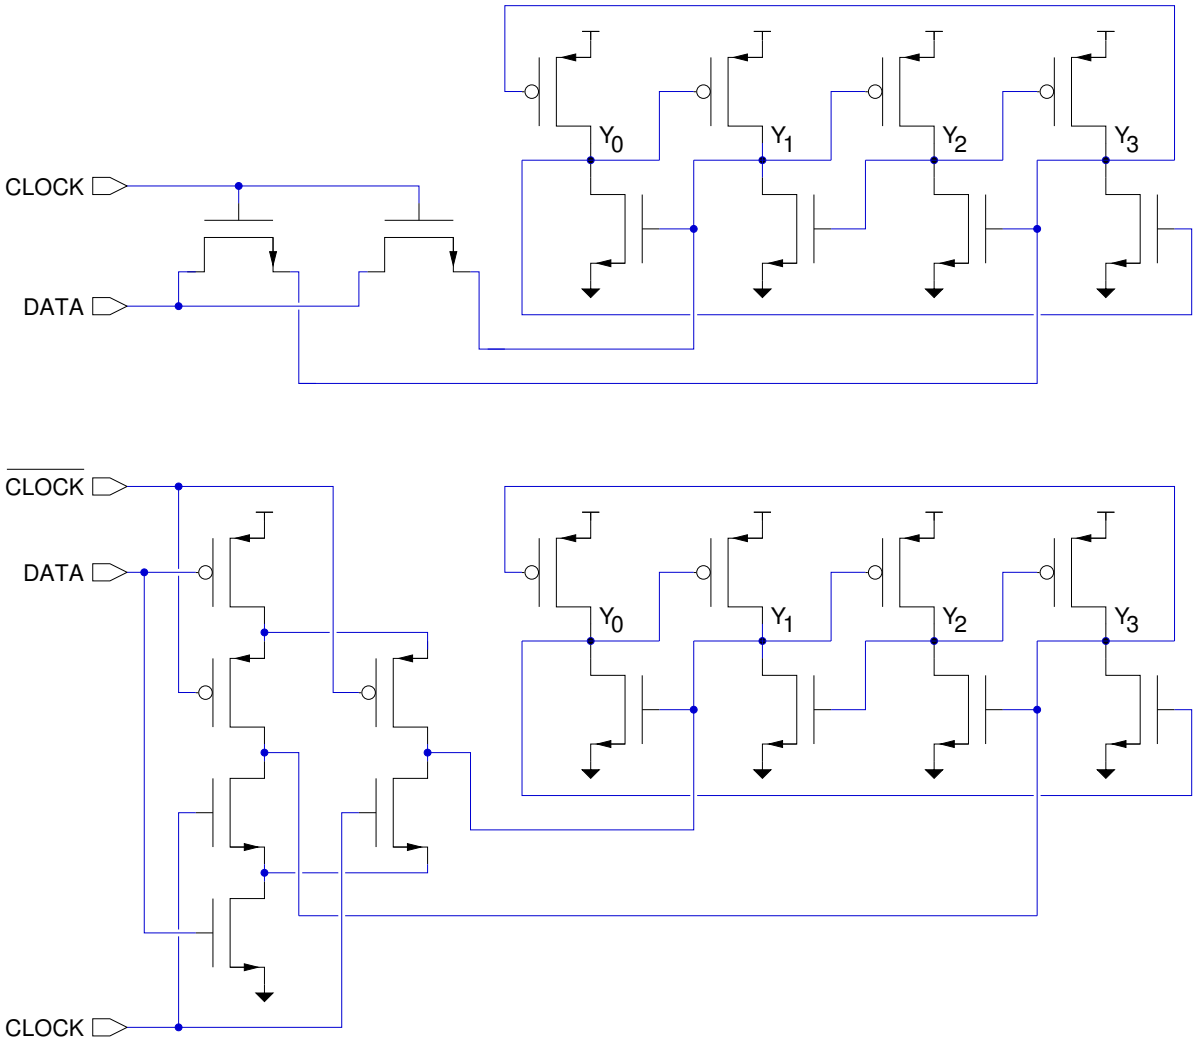

Figure 4.5: DICE input circuits

required to predischarge the cell. The second situation is where a new data value is magnetically written into the MTJs and the latch is required to assume the new logical value immediately after the write operation. For the circuits described here, where the MTJs are connected between the write line and the NMOS bottom gates, the voltage applied to the write line will enable the NMOS transistors, causing them to sink current and discharge the cell. The final application scenario applies when the state of the MTJs must be resolved at some arbitrary time, completely independent of any power-up sequence or MTJ write operation. It becomes necessary in this case to add circuitry to the latch to insure a complete discharge of the internal nodes, which can be accomplished simply by adding NMOS transistors in parallel

with Q5N through Q8N that have their front gates controlled by the LDN signal.

4.3 Design Issues

In order to reveal the critical design issues in the shadow latch circuit it is helpful to consider the behavior of the transistors alone. Suppose that the MTJs are removed from the circuit of Fig. 4.2 and the bottom gates of the NMOS transistors are tied to ground. Assume that the transistors are perfectly matched, all parasitic circuit elements can be neglected, and all internal nodes are at ground: the cell is perfectly balanced. When LDN is brought low to enable Q5 the PMOS transistors Q3 and Q4 will source current into DATA and DATAN because the potential at their sources has been brought much higher than the potential at their gates and drains. The voltages at DATA and DATAN will continue to increase simultaneously and equally, so that the V_{GS} of the PMOS transistors decreases as the V_{GS} of the NMOS transistors increases. At some point, DATA and DATAN will reach a voltage such that $I_{DSP} = I_{DSN}$ and the circuit will enter a metastable state. In an ideal world (or a SPICE simulation) this state could persist indefinitely, but the inverters have very high gain under these conditions so any noise coupled into the circuit will cause it to flip into a truly stable state with all nodes at either V_{DD} or V_{SS}. However, the more likely scenario is that the transistors will not be perfectly matched and the parasitic elements will not be negligible, making the cell unbalanced. In this case either DATA or DATAN will rise slightly faster than the other, causing to cell to flip to a preferred logical state.

The MTJs are added to the circuit with the intent of giving it a programmable imbalance. Since the difference between the programmed states of the MTJ is relatively small, the imbalance that they impart to the cell is also small. Therefore, the first critical design issue is that all other sources of imbalance in the cell must be minimized. Designers of analog circuitry are very familiar with techniques for matching transistors and reducing parasitic

effects. Some of these can be applied to the latch layout and are summarized in Table 4.1.

Source	Mitigation
gate length variation	gate lengths greater than minimum
gate width variation	gate widths greater than minimum
implant shadowing	same orientation for all channels
source/drain resistance	source and drains covered with Metal0
contact/via resistance	redundant contacts and vias
metal resistance	metal width greater than minimum
wiring capacitance	equalize wire lengths

Table 4.1: Sources of transistor imbalance

An additional source of inherent imbalance is any variation in the voltages at DATA and DATAN when they are expected to be at ground (at the instant when LDN is asserted). In the circuits of Fig. 4.2 and Fig. 4.4, the "grounded" end of the MTJs will be brought to a higher potential during the process of programming the MTJs, which will enable the NMOS transistors and discharge the internal nodes. An important design parameter is thus to determine the minimum time required under worst case conditions for the nodes to completely discharge. If the discharge time is excessive, additional NMOS transistors, enabled by LDN, can be added in parallel with the existing NMOS transistors.

Having reduced the sources of unwanted imbalance in the cell, the next step is to maximize the desired imbalance produced by the MTJs. The goal is to imbalance the cell by altering the threshold voltage of the NMOS transistors, and this is accomplished by varying the voltage applied to the bottom gate. The imbalance should be created during the short period when the internal node voltages are rising after LDN is asserted low but before the metastable state is reached. Once the cell enters the metastable state it is vulnerable to switching by any random physical imbalance or induced noise, so the MTJs should introduce their imbalance earlier. Another consideration is the small signal gain from the bottom gate to the drain as a function of the d.c. bias at the bottom gate. The simple circuit shown on the left of Fig. 4.6 was extracted from the complete DICE cell and simulated to produce the results given in the graph on the right. The voltage applied to the bottom gate V_{BG}, comprised an a.c. signal of 5 mV at

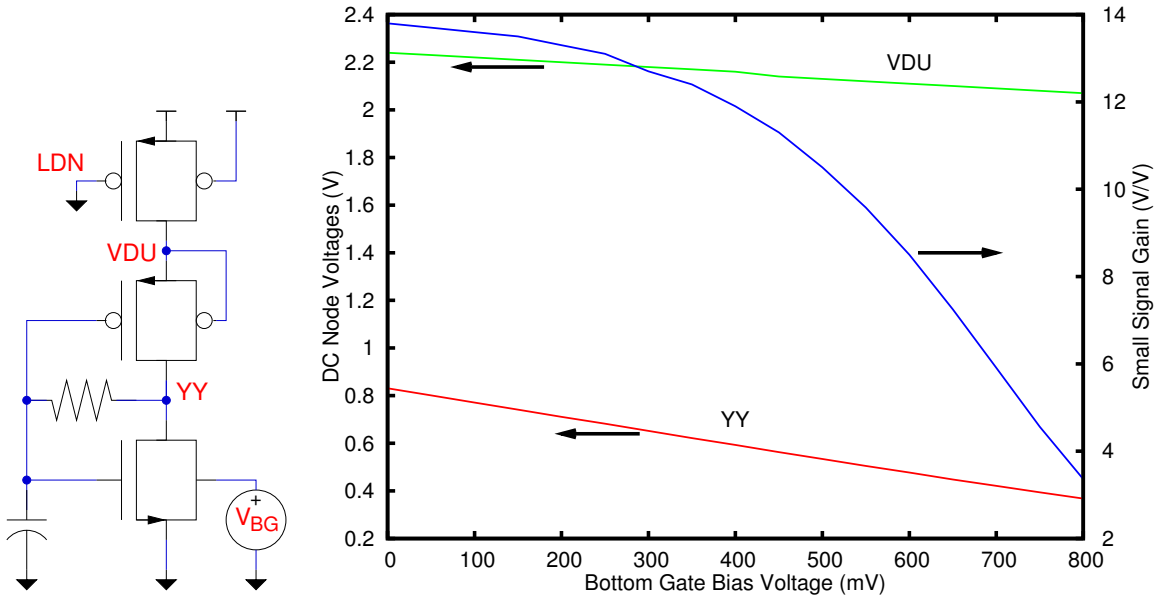

Figure 4.6: Bottom gate small signal gain

100 MHz and a d.c. bias voltage, shown as the *X* coordinate of the graph. The resulting d.c. voltages at VDU and YY are plotted against the *Y* axis on the left, while the small signal voltage gain from the bottom gate to YY is plotted against the *Y* axis on the right. Note that the d.c. voltage at YY is relatively low in the metastable state, well below $V_{DD}/2$, because the ability of the cell to recover from a cosmic particle requires that the NMOS transistors have significantly greater drive strength than the PMOS. As the bottom gate bias increases the front gate threshold voltage decreases, increasing the drain current through all of the transistors and pulling YY low. Unfortunately, this also reduces the gain of the inverter so any circuit modifications that increase the absolute magnitude of the bottom gate voltages rather than ΔV may be counterproductive.

In addition to these functional requirements, the secondary goals of low power consumption and high reliability should be addressed. To reduce the leakage current through the NMOS transistors and eliminate any d.c. current through the MTJs, the bias across them should be reduced to zero after the cell reaches a stable state. During switching, the MTJ bias should be kept well under 1 V to prevent damage to the tunneling barrier.

The circuits shown in Fig. 4.2 and Fig. 4.4 induce a voltage on the NMOS bottom gates with capacitive coupling to the transistor's top gate (and drain, to a lesser degree). The transistor gate dimensions and the bottom gate wiring parasitics should be designed to produce the largest induced voltage on the bottom gate, without forward biasing the junction from the bottom gate to the source. At the same time, the RC time constant of the MTJ and bottom gate must be maintained within appropriate bounds. If the time constant is too small the MTJs will completely discharge the bottom gates, eliminating any imbalance, before the cell can flip to the desired state. If the time constant is too large then the imbalance developed during the switching time will not be sufficient to overcome any inherent imbalance in the cell.

4.4 Design Optimization

The design shown in Fig. 4.4 was used as a starting point for the design process. Two design tools from SIMUCAD Design Automation, *Expert* and *HIPEX* were used to capture the circuit layout and extract netlists, respectively. Circuit simulations were performed using Synopsys' *HSPICE* simulator. Layout design rules, process parameters, and transistor models were provided by American Semiconductor and this work used version 12 of their design kit, dated 24 August 2006. Unless explicitly stated otherwise, timing measurements are made at a voltage level equal to one-half of the supply voltage, which is nominally 2.5 V. The initial W/L ratio for transistors `Q5N` through `Q8N` and `Q5P` through `Q8P` was 0.9 μm/0.18 μm, and transistor `Q0` was drawn as 3.0 μm/0.18 μm.

4.4.1 Read Operation With Grounded WRTLN

To simplify the evaluation and eliminate extraneous variables, the first simulations of the circuit were made solely to evaluate its ability to discern small magnetoresistance ratios. Therefore, the starting point for the simulations was with the common node for the MTJs, the

write line, at ground. All internal circuit nodes were also initialized to ground. The MTJs were replaced with ideal resistors, one which had a fixed, nominal value, and the other which was varied around that nominal value. The simulations were repeated for nominal resistance values from 15 kΩ to 35 kΩ in steps of 5 kΩ. The other resistor was varied over a range of ±30% from this nominal value in steps of 1%, and the time required for the latch to resolve a '0' or '1' state after the falling edge of LDN was measured. The results are shown in Fig. 4.7 for four different circuit netlists. The lines plotted in green, on the left side of each graph, are for cases where the latch resolved to a '0' state while the red lines on the right side are for cases where the latch resolved to a '1' state. For small resistance ratios the circuit required more time to resolve one state or the other, and in some cases it entered a metastable state with no resolution within 5 ns.

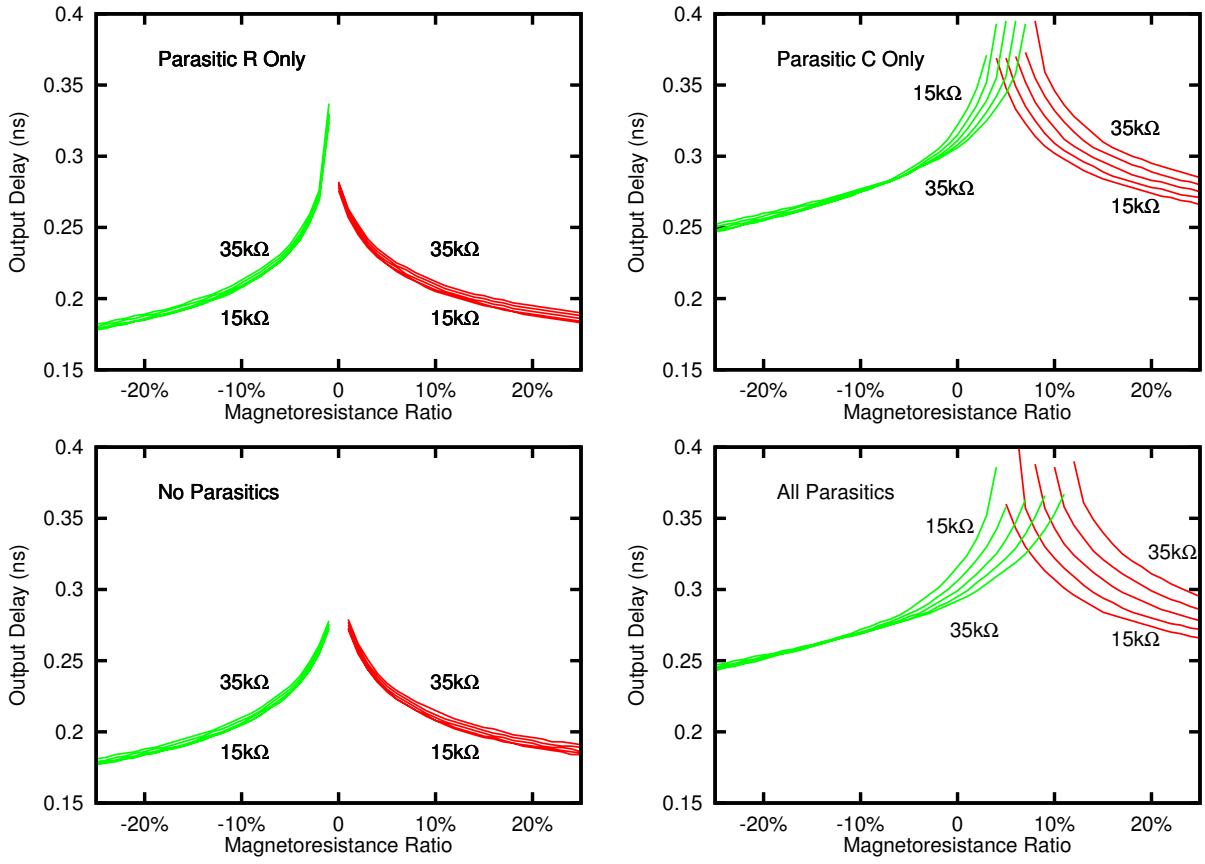

Figure 4.7: Shadow RAM latch simulations replacing MTJs with resistors

The lower left graph of Fig. 4.7 shows the circuit behavior without any parasitic wiring resistance or capacitance. In other words, the netlist contained only transistors. As expected, the output delay values are centered and symmetric about the 0% magnetoresistance ratio (i.e. the case where both resistors had the same value). As the nominal resistance value increased from 15 kΩ to 35 kΩ the output delay increased slightly. The upper left graph shows the results when the netlist included 983 resistors representing the wiring resistance. The output delays are only slightly higher than the first graph, because the wiring resistances in this small circuit are not large. In the upper right graph are the results when the netlist included the capacitance (as 86 lumped capacitors), but not the resistance, of the wiring. In this case the output delays vary significantly as the nominal MTJ resistance increases, and reveal an offset in the response curves. Even when the resistance ratio is positive, so that a '1' should be resolved, the circuit may declare the stored value to be a '0'. For a nominal resistance of 15 kΩ the threshold between logic states lies at a resistance ratio between +3% and +4%. When the nominal resistance is 35 kΩ this threshold moves to the right on the graph, between 7% and 8%. Finally, on the lower right, both the parasitic wiring capacitance (2699 capacitors) and resistance (983 resistors) are included, and the resolution threshold has increased to something between 11% and 12% for the highest nominal resistance value.

Clearly, the parasitic capacitances in the layout are introducing an undesirable skew in the latch's ability to resolve small magnetoresistance values. The relevant capacitances for the MTJ nodes that are connected to NMOS bottom gates are shown in Table 4.2. The most

	BG13	BG02
WRTLN	0.28 fF	0.36 fF
VSS	0.30 fF	0.32 fF
Substrate	1.11 fF	1.05 fF

Table 4.2: Parasitic capacitances in first shadow latch

striking difference is the capacitance to WRTLN, the write line. This is explained by examining the layout in Fig. 4.4, where the wiring around the two MTJs is slightly different.

The connection to the bottom of the MTJ on the left is node BG13, and it uses less Metal1 wiring and more Metal0 wiring than the connection to the bottom of the MTJ on the right, node BG02. Since the write line passes over the MTJs in Metal2, BG13 will have less capacitance to WRTLN, and the additional Metal0 on this node will increase its capacitance to the substrate. (For this SOI process, the "substrate" is the conductive part of the wafer that is buried beneath the oxide layer, so it is distinct from the transistor body terminals and is only coupled to the circuit capacitively. For simulation stability purposes the substrate is connected to VSS with a 10 MΩ resistor.) Adding a small ideal capacitance between WRTLN and BG13 confirmed that this difference was the primary source of the skew in the latch.

A new layout of the latch, shown in Fig. 4.8 was created with the intent of equalizing the

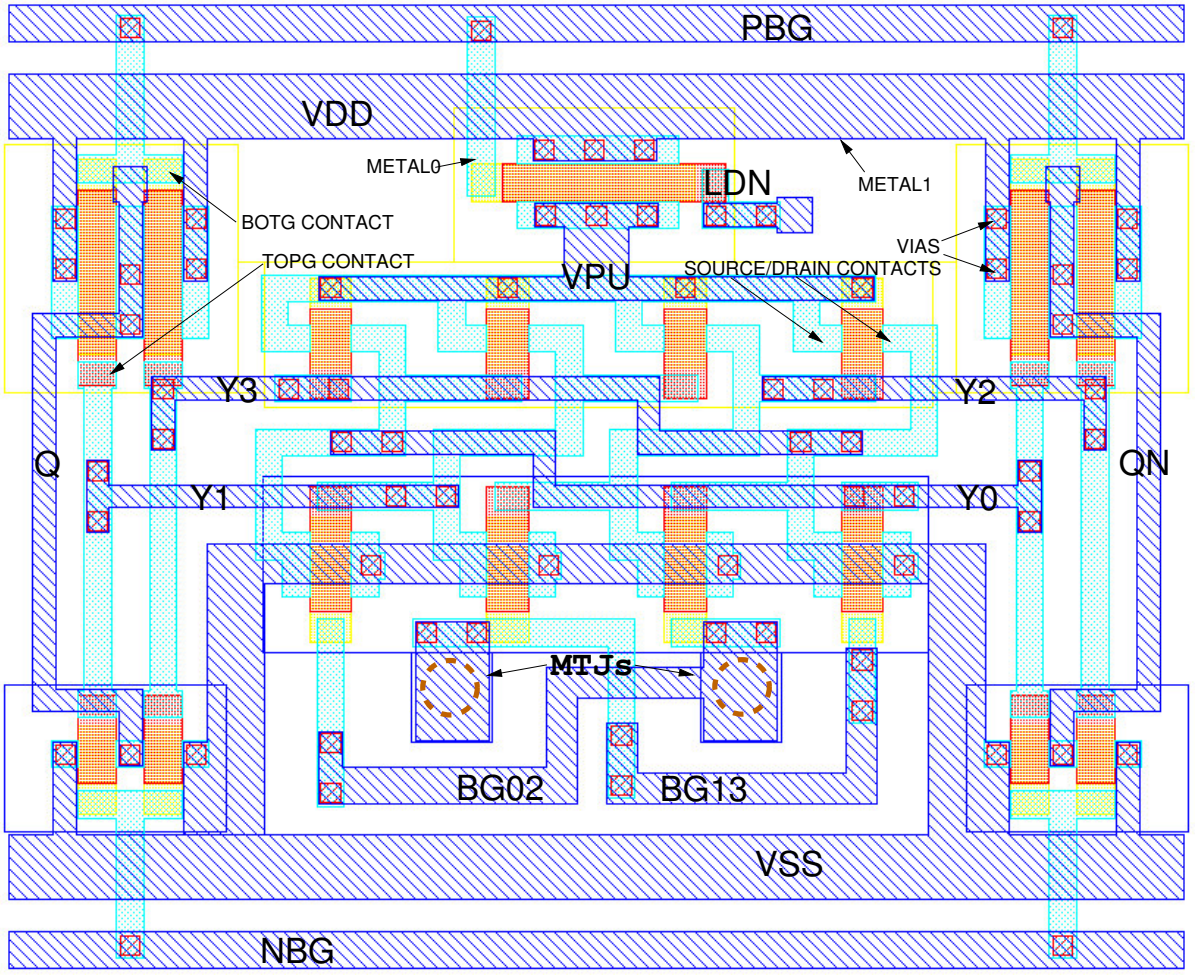

Figure 4.8: Shadow RAM cell layout, second version

parasitic capacitances at the MTJ nodes. (The SPICE netlist is provided in Appendix E.) In this layout the capacitance from WRTLN to BG13 increased to 0.37 fF while the capacitance from WRTLN to BG02 increased to 0.40 fF. The capacitances from BG13 and BG02 to VSS decreased to 0.12 fF and 0.15 fF respectively, and the capacitances to the substrate were reduced to 0.92 fF and 0.96 fF.

At the same time, some transistor lengths and widths were increased to reduce imbalances caused by random lithographic variation. The drawn gate length of transistors `Q5N` through `Q8N` and `Q5P` through `Q8P` was increased by 50% to 0.24 μm. In order to insure that the NMOS transistors in the DICE latch would always overdrive the PMOS transistors, and to increase the capacitive coupling from the top gate to the MTJs, the width of `Q5N` through `Q8N` was increased to 1.47 μm. The circuit simulations of this layout with no parasitic resistance or capacitance, and with full extraction of parasitics (970 resistors and 2315 capacitors), are shown in Fig. 4.9. As desired, the skew in resolution of '0' and '1' states has been greatly

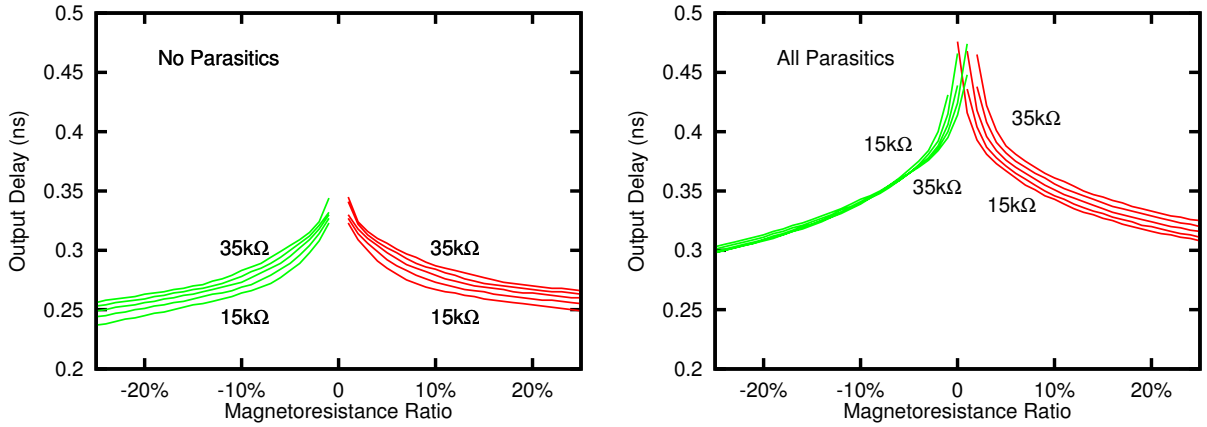

Figure 4.9: Shadow RAM latch simulations replacing MTJs with resistors, second layout

reduced, but the larger transistor gate lengths have generally increased the circuit delay.

The graph in Fig. 4.10 shows detailed simulation results for the the behavior of this circuit layout. This simulation was performed using the MTJ model described in Appendix B, with a resistance of 15 kΩ for the parallel MTJ and a TMRR of 30% at zero bias. The capacitance of the MTJ was included, with a relative permittivity of 8.0 (Al_2O_3), a dielectric thickness of

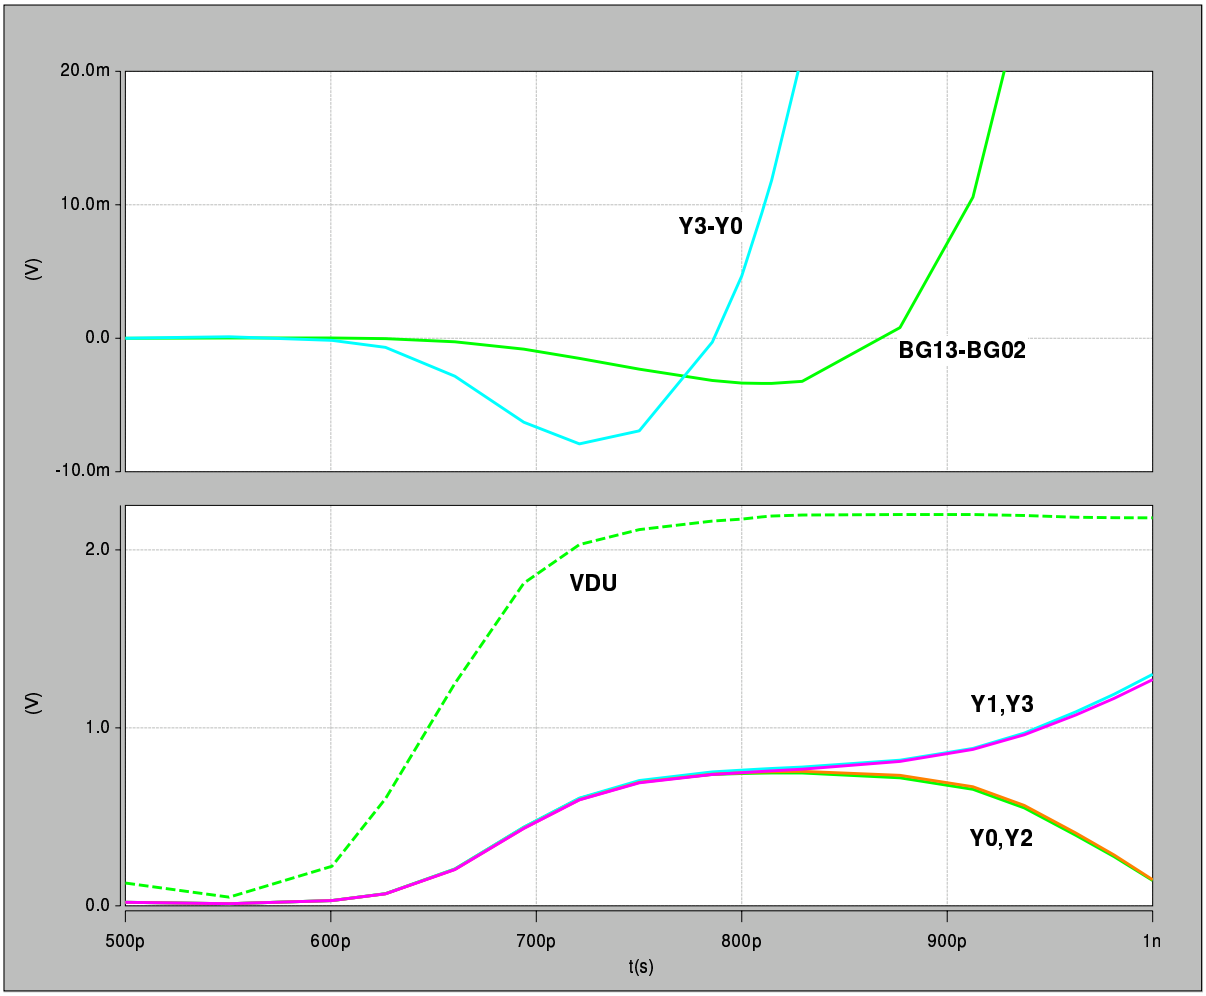

Figure 4.10: Simulation detail, second layout

3.5 nm, and an area of 0.5 μm^2. The resulting capacitance of the junction was 10.1 fF. By comparison, the total parasitic wiring capacitance was 1.64 fF on BG13 and 1.68 fF on BG02, with only a small fraction of the total being coupling capacitance to Y0 through Y3. `MTJ1`, connected to BG13, was assumed to be in a parallel orientation while `MTJ2` was antiparallel.

The critical behavior of the circuit as it resolves the relative resistance of the two MTJs can be seen at around 800 ps in Fig. 4.10. The bottom set of traces shows that VDU has reached a relatively stable voltage, and all four of the DICE cell nodes (Y0 through Y3) are at the same voltage in a metastable condition. The top two traces show the voltage difference between two node pairs. As the cell moves toward the metastable condition there is a small difference

between Y3 and Y0 due to differences in parasitic capacitance and resistance at these nodes, but the offset is only about 1% of the node voltages. More importantly, the TMRR of the MTJs is reflected in the voltage difference (or ΔV) between the NMOS bottom gate connections, shown in the trace of BG13-BG02. Once a ΔV of several millivolts is established the latch begins to leave the metastable condition, with Y1 and Y3 climbing toward V_{DD} and Y0 and Y2 falling to ground. Although not shown in Fig. 4.10, if `MTJ1` is antiparallel and `MTJ2` is parallel, the shape of the BG13-BG02 trace is virtually unchanged except that the polarity is inverted (the trace is mirrored about the *x* axis), with the corresponding reversal of the traces for Y0 through Y3.

The HSPICE simulations show that the shadow latch circuit can correctly resolve the state of the MTJs, but the observed ΔV for the bottom gate voltages is disconcertingly small. At this point it is reasonable to try to optimize the MTJ design parameters in order to maximize the obtainable ΔV and minimize the time required for the latch to produce the correct output. The parameters available for "tweaking" are primarily the resistance-area (RA) product of the junction and its capacitance. Since the capacitance of the junction varies in proportion to its area and the resistance is inversely proportional to area, the RC time constant inherent in the junction is more or less independent of the junction area and is determined by RA, the insulator thickness, and the relative permittivity of the dielectric.

It seems intuitively clear that reducing the capacitance of the junction should result in higher bottom gate voltages, because the junction capacitance forms a capacitive voltage divider with the capacitance from the top gate and drain. On the other hand, increasing the bias across the MTJ will decrease the TMRR, so the relationship between the peak bottom gate voltage and ΔV will not be monotonic and for bottom gate voltages above about 500 mV ΔV falls to zero. Having determined the junction capacitance, the desired RA product will have a value that results in a t_{RC} comparable to the time required for the latch to enter the metastable state, where the inverter gain is highest. Fig. 4.11 presents simulation results for various values of

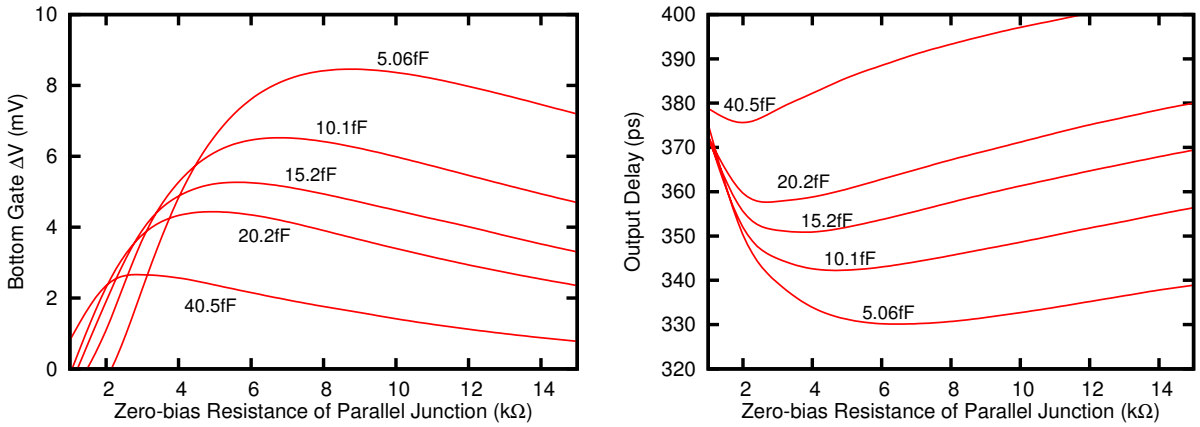

Figure 4.11: MTJ optimization, second layout

the nominal resistance (RA times area) and capacitance of the MJTs, with the TMRR set at 30%. The left graph in the figure indicates how the difference in voltage between the two bottom gate terminals, at one terminal of each MTJ, varies as a function of the junction resistance. The voltage difference between the bottom gates is recorded at the instant in time when the difference between the DICE nodes Y0 and Y3 reaches 50 mV, which has been (somewhat arbitrarily) selected as the "decision point", the point where the latch has left the metastable condition and has resolved the state of the MTJs. In the right graph is plotted the time delay from the falling edge of LDN to the falling edge of Q when the latch resolves to a '0' state. Over a reasonable range of junction capacitance it appears that a nominal resistance of approximately 5 kΩ results in the largest ΔV and fastest output delay. As expected, resistance values that are much larger or much smaller than this optimum range will sacrifice the performance of the latch. Note that unlike the differential latch described in Section 3, a high resistance is not needed to provide SEU immunity in the shadow latch. Another conclusion that can be drawn from these graphs is that there is a strong correlation between the magnitude of the ΔV induced on the NMOS bottom gates and the speed with which the latch resolves to the desired state.

In the layouts and simulations described above, the active capacitive coupling from the NMOS bottom gates is primarily to the NMOS top gates. The channel acts as the insulator

between these nodes, resulting in a net capacitance of less than 1 fF from top to bottom. An obvious approach for increasing ΔV is to increase the magnitude of the voltage coupled to the bottom gates by increasing the capacitance from them to a chosen circuit node. There are three reasonable candidates for the "chosen circuit node" in this case. One could simply increase the capacitance from top gate to bottom gate, add parasitic capacitance from the bottom gate to the transistor drain, or add capacitance from the bottom gate to node VDU. These options were simulated for two values of parallel MTJ resistance, 5 kΩ and 15 kΩ, with a TMRR of 30% and an MTJ capacitance of 10.1 fF. The results of these simulations are shown in Fig. 4.12.

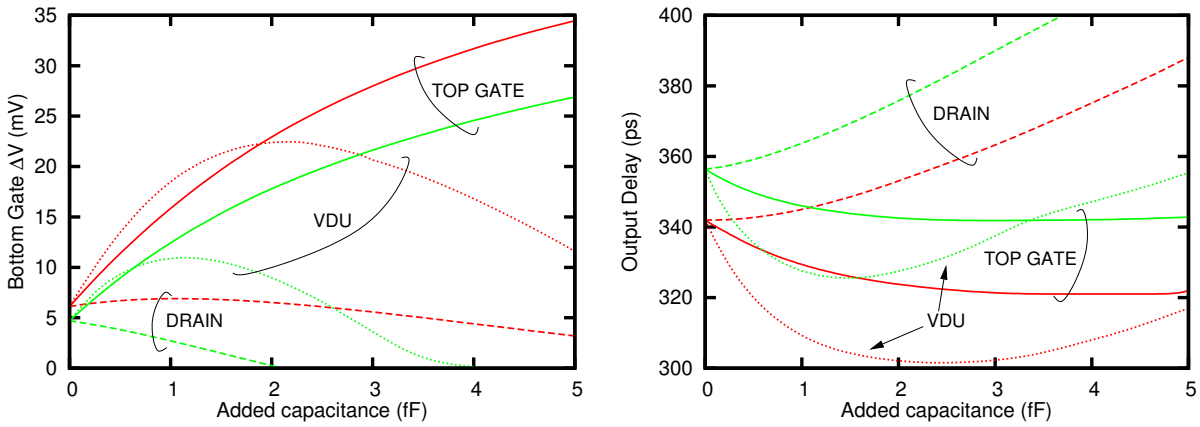

Figure 4.12: Effect of added capacitance, SPICE data

As before, the graph on the left indicates the magnitude of ΔV when the latch reaches its decision point. There are three pairs of curves, corresponding to a fixed capacitance added between the bottom gate and either the top gate, the transistor drain, or VDU. In each pair of curves on the left the upper (red) curve is for a 5 kΩ nominal resistance and the lower (green) curve is for the 15 kΩ resistance. The graph on the right indicates the output delay time for the three possible capacitance connections, but in this case the upper (green) curve is for the 15 kΩ nominal resistance and the lower (red) curve is for the 5 kΩ resistance. Clearly, adding capacitance between the bottom gate and drain is a poor choice. Coupling to the drain actually creates negative feedback that tends to eliminate the desired ΔV, because the drain node of

each NMOS transistor is moving in the opposite direction as the desired change in the bottom gate. For example, the voltage at the bottom gate connected to the parallel MTJ should fall more quickly, which means that the transistor's threshold voltage is increasing more quickly, which means that the transistor's drain voltage will be rising. On the other hand, increased coupling to the top gate provides positive feedback and reinforces the resistance difference of the MTJs. As the capacitance to the front gate increases the ΔV also increases and the output delay decreases, although the improvement in output delay tends to saturate for capacitances greater than about 3 fF. Coupling the bottom gate to VDU provides no feedback, but benefits from the faster edge rate and earlier occurrence of VDU. However, as this capacitance is increased beyond some optimal point the RC time constant of the circuit causes a delay in reaching the peak ΔV, and the benefit is lost.

Note from Fig. 4.12 that, for a capacitance of 2 fF, coupling to either the front gate or VDU produces the same ΔV but the circuit with coupling to VDU is significantly faster. The reason for this difference can be seen in the simulation waveforms shown in Fig. 4.13. In these plots the solid traces represent the behavior of the original circuit, the dashed traces are for the circuit with 2 fF added between the top and bottom gate of NMOS transistors `Q5N` through `Q8N`, and the dotted traces are for 2 fF added between the bottom gates of those transistors and VDU. The dV/dt on VDU is much higher than that of the nodes connected to the top gates of `Q5N` through `Q8N` (Y0 through Y3). Consequently, the bottom gate voltages are much higher, as shown in the top graph of Fig. 4.13. Note that with added capacitance to VDU the bottom gate nodes reach about 250 mV, twice the peak voltage with added capacitance to the top gates, and roughly five times the peak voltage with no added capacitance. The higher bottom gate voltage produces lower NMOS thresholds, increased drive current, and faster circuit resolution. These curves also agree with the data presented in Fig. 4.12 in that the ΔV between BG02 and BG13 is roughly the same (about 23 mV) at the decision point (about 800 ps, where Y3-Y0 reaches 50 mV) for both cases of added capacitance. However, the ΔV for capacitance to VDU begins to decay after this point while the positive feedback provided

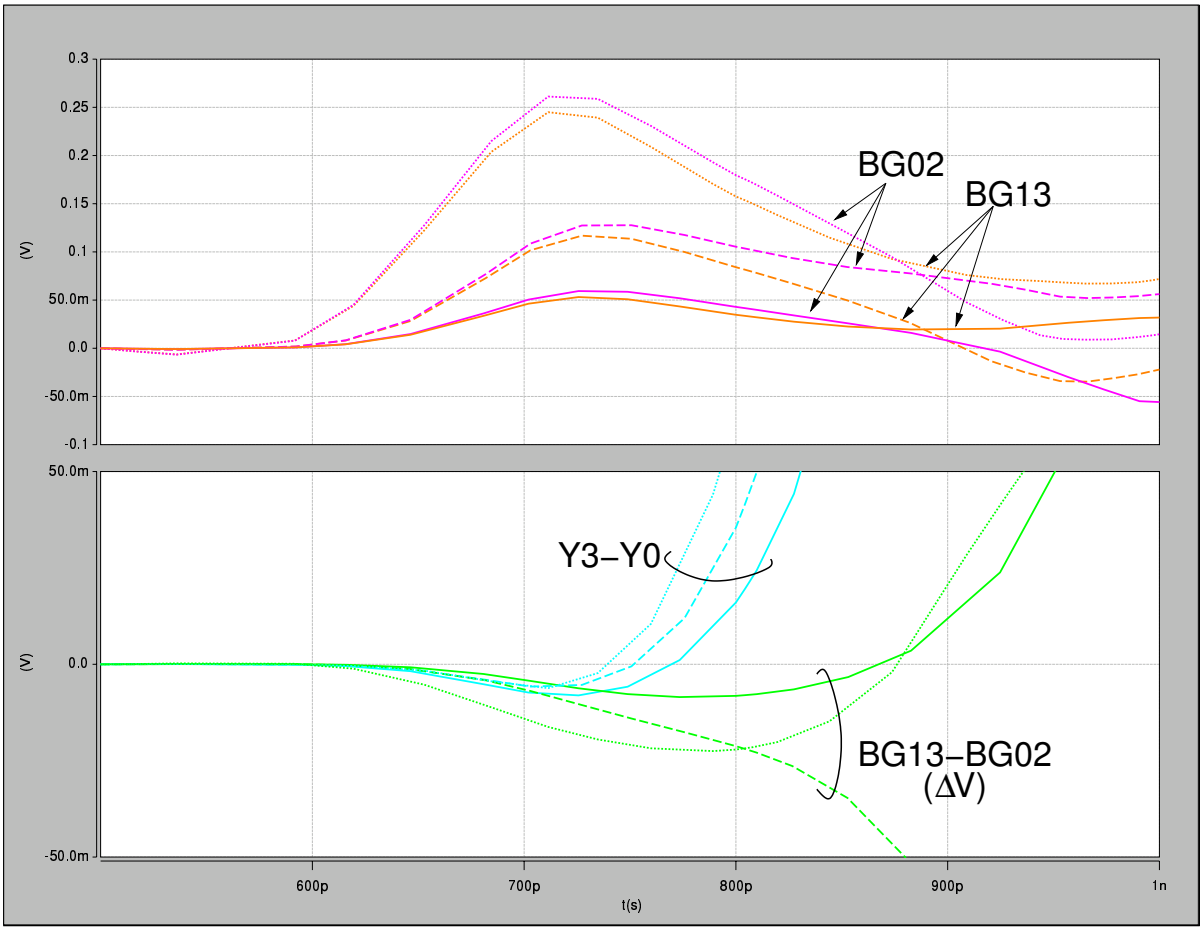

Figure 4.13: Effect of added capacitance, waveforms

by capacitance to the top gate causes ΔV to continue to grow.

Unfortunately, a 2 fF capacitance can be a relatively large structure in a sub-micron SOI technology. Implementing this using the parasitic capacitance between interconnect layers would be prohibitively large, so the only viable option is to use the MOSFET gate-to-channel capacitance. American Semiconductor recommends using a FlexFET with the top gate as one terminal and the source, drain, and bottom gate connected together as the other terminal of the capacitor. An NMOS FlexFET configured in this manner with a width of 0.9 µm and a gate length of 0.24 µm provides approximately 2 fF and fits reasonably well in the shadow latch layout, as shown in Fig. 4.14. (The SPICE netlist is provided in Appendix F.) The added capacitors are shown in the schematic as transistors `Q5C` through `Q8C`, located just below

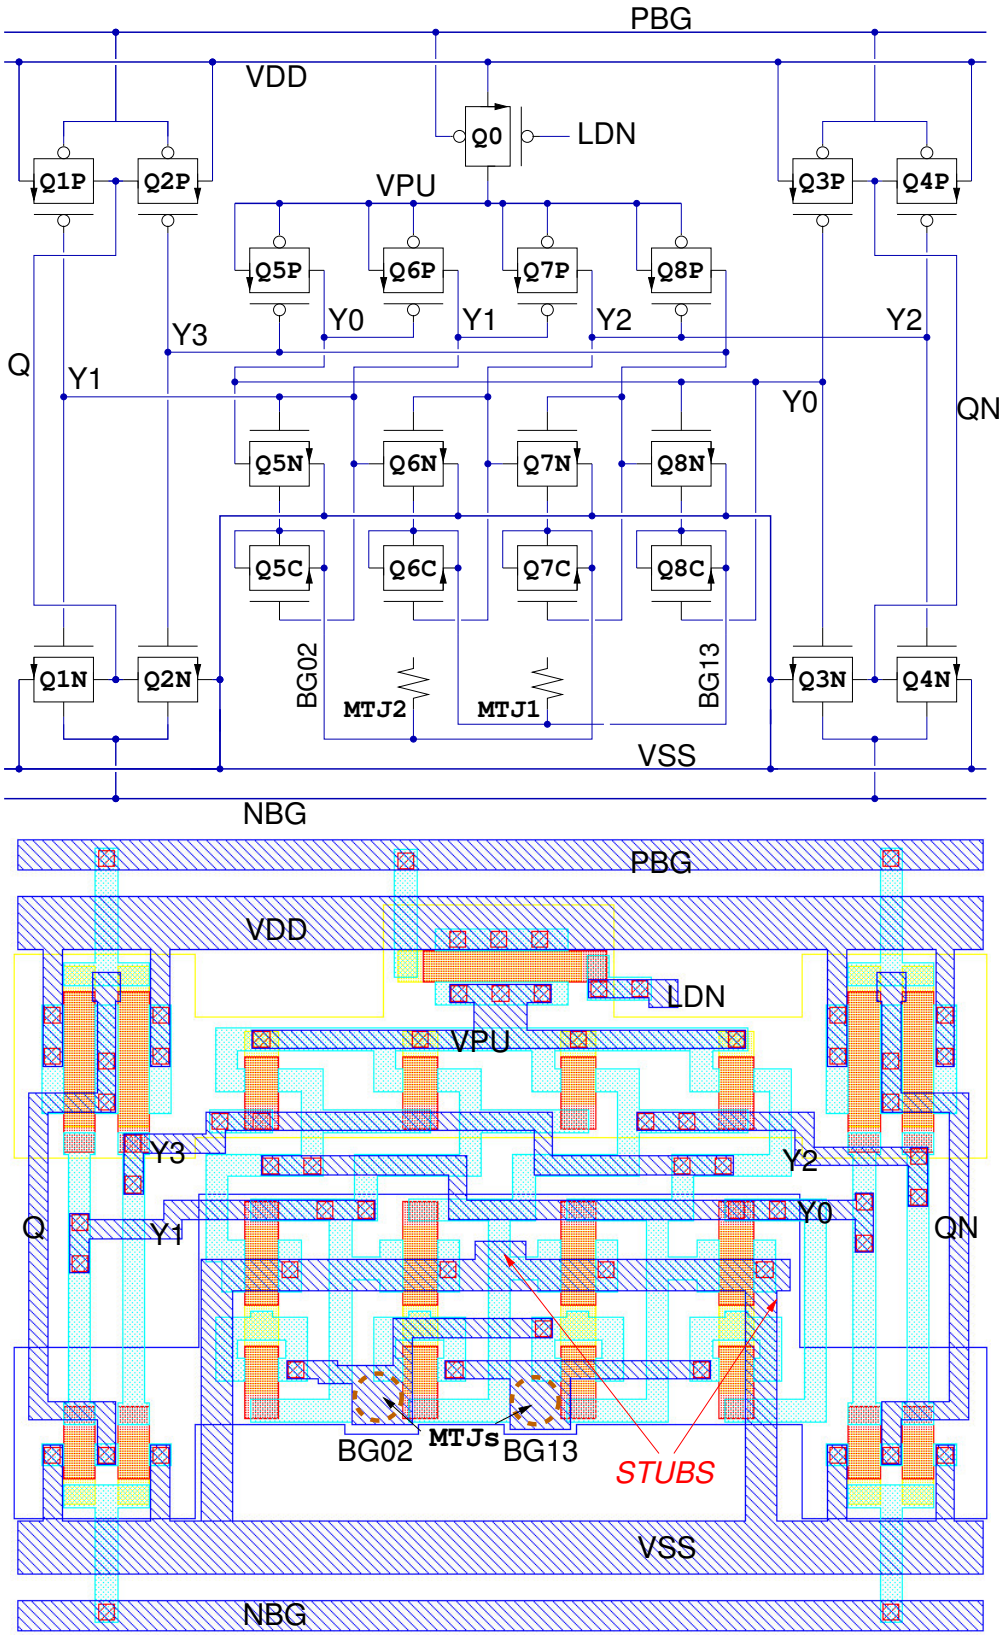

Figure 4.14: Shadow RAM cell layout, third version

transistors Q5N through Q8N. This arrangement also exists in the layout, where the capacitors are oriented so that their bottom gate terminals are toward the top of the page so that they connect by abutment with the bottom gate terminals of the original NMOS transistors. The top gate connection is brought down on the right side of each capacitor. The connections to the MTJs were rearranged to maintain equal parasitic capacitance, with the result that MTJ1, connected to BG13, is now on the right and MTJ2, connected to BG02, is on the left.

Note the two stubs in Metal1 that are indicated on the layout. These were added solely to balance the capacitance from VSS to Y0 through Y3. Simulations revealed that equalizing the parasitic capacitances on Y0 through Y3, as well as the capacitance on BG02 and BG13, is critical to the ability of the latch to resolve small TMRR values. In general, all parasitic capacitances from any given node to all other nodes that are static, d.c. potentials (such as VDD, VSS, NBG, and PBG) can be lumped together for this purpose. However, for the FlexFET SOI process the largest single parasitic capacitance for these nodes is to the buried substrate, which may or may not fit into this category. For the substrate to be considered a static node it must be connected to one of the power supply terminals with a low impedance, such as through a down-bond and conductive die attach to the floor of the package cavity.

The relative importance of matching the various parasitic capacitances is illustrated in the *pairs plot* diagram in Fig. 4.15. This diagram reveals the correlations between pairs of simulation variables and a chosen measured behavior. In this case the simulation variables are the *mismatch* in parasitic capacitances from BG02 and BG13 to some other circuit node. The data shown resulted from an HSPICE Monte Carlo simulation, where 2500 random combinations of five mismatch variables were evaluated. For each mismatch variable, the capacitance from BG13 to the chosen node was fixed at a nominal value and the capacitance from BG02 to the same node was randomly and uniformly varied ±30% from that nominal value. Therefore, the X and Y axes of each plot represent a range of 0.7 to 1.3 times the nominal capacitance to two different nodes.

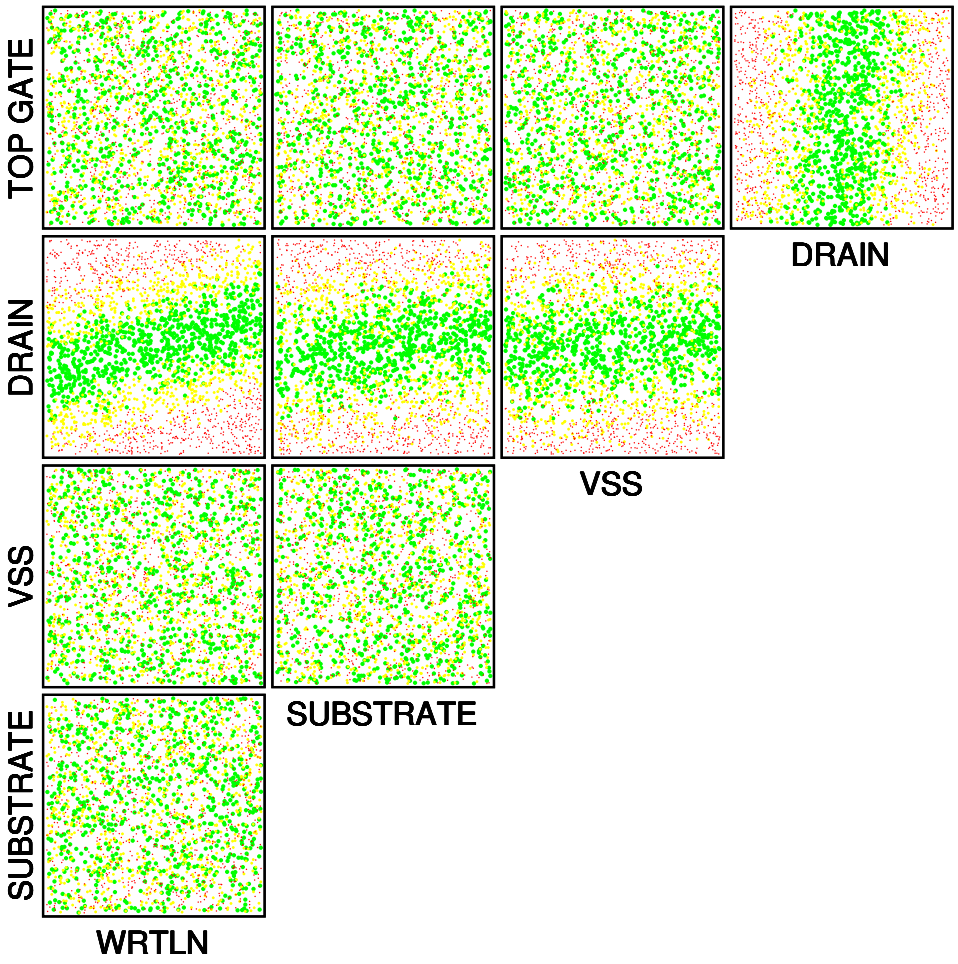

Figure 4.15: Pairs Plot, capacitance mismatch and delay skew

The measured behavior in these simulations is the difference in the propagation delay for the two logic states of the latch. As was shown earlier, this delay skew is a reflection of the circuit's ability to resolve small magnetoresistance ratios. Each dot in a graph represents a simulation that was performed, with its location in the graph corresponding to the skewed values of the two mismatch variables being evaluated and the color of the dot corresponding to the observed delay skew. Green dots are plotted when the skew was less than 5 ps, yellow dots are plotted for skew values between 5 ps and 10 ps, and red dots are for skew values greater than 10 ps. Each row of graphs in Fig. 4.15 has the same variable used as the Y axis, and each column has the same variable used as the X axis. The arrangement of small graphs shown in the figure allows the correlation between all possible pairings of the five mismatch

variables to be visualized. For these simulations a TMRR value of 30% and nominal resistance of 5 kΩ were used.

For example, consider the bottom left graph in the Fig. 4.15. The *Y* axis represents the difference, or mismatch, between the capacitance from BG13 to the substrate and the capacitance from BG02 and the substrate. Similarly, the *X* axis represents the capacitance mismatch from BG13/BG02 to the WRTLN node. The random pattern of colored dots in this graph suggest that there is little correlation between either of these capacitance mismatches and the output delay skew. However, all of the graphs shown in the row labeled DRAIN have a distinct grouping of green dots. The green dots tend to form a horizontal band, indicating a strong correlation between the mismatch in capacitance to the NMOS transistor drain terminals and the resulting delay skew. Since the *Y* axis range is from -30% mismatch to +30% mismatch, the location of the green band near the middle of the *Y* axis indicates that the optimum mismatch for this capacitance is zero. Note that the left-most graph in the DRAIN row shows a clear positive slope, indicating a small positive correlation between the capacitance mismatch to the top gate nodes, the capacitance mismatch to WRTLN, and the delay skew. The implications for circuit design are clear: if the capacitance from one of the MTJs to its associated drain nodes is less than that of the other MTJ, we can somewhat compensate for this by reducing the capacitance from that MTJ to the WRTLN.

Note that the effect of random variation in the parasitic capacitance to the top gate nodes has little effect on the output delay skew, because this parasitic capacitance is quite small when compared to the added capacitance of Q5C through Q8C. Capacitance mismatches to the relatively static nodes (WRTLN, VSS and the substrate) are also relatively small in magnitude and there is no feedback mechanism from these nodes that would amplify the effect of a small variation.

At this point some general observations about the pairs plot graphs can be made. A horizontal line or band indicates a strong correlation between the measured result and the variable

plotted on the *Y* axis, a vertical band indicates a strong correlation to the variable on the *X* axis, and a small cluster of green dots would indicate a strong correlation between both input variables and the measured result, but no correlation between the input variables. A band with a positive slope indicates a positive correlation between the two input variables while a negative slope shows a negative correlation. The power of the pairs plot is that it allows a great deal of information to be extracted from a body of random simulations.

The relationship between the capacitance mismatches can be examined in a more traditional manner by sweeping the value of one mismatch while measuring the output delay skew, and repeating the simulation for several values of the other mismatch, as shown in Fig. 4.16. The capacitance from BG13 to WRTLN was fixed at 0.85 fF while the capacitance from BG02 to

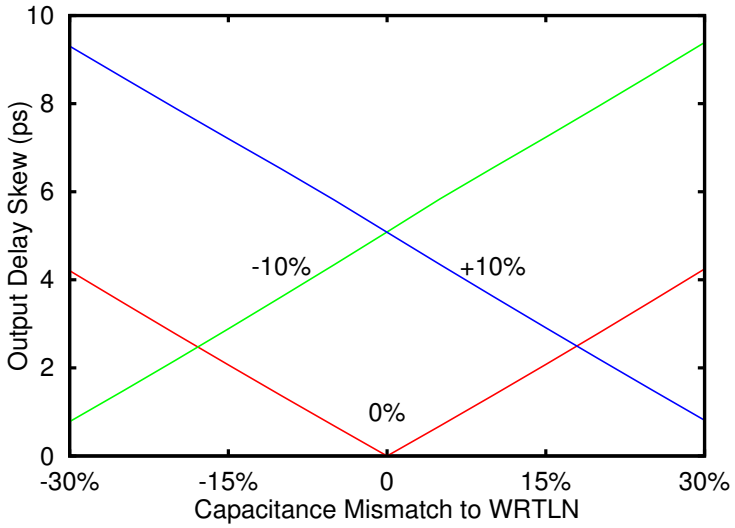

Figure 4.16: Output delay skew vs. capacitance mismatch to WRTLN and capacitance mismatch to drains

WRTLN was varied ±30% from this value. Similarly, the capacitance from BG13 to both Y1 and Y3 was fixed at 0.1 fF and the simulation was repeated for capacitances from BG02 to Y0 and Y2 that were equal to 0.1 fF and ±10%. The resulting magnitude of skew between the 0 and 1 state output delays is plotted in Fig. 4.16. As expected, when both capacitance mismatch values are at zero there is no skew in the output delays. On the left side of the graph we can see that a -30% mismatch in the capacitance to WRTLN alone causes a delay skew of

just over 4 ps, but a simultaneous mismatch of -10% in capacitance to the drains will reduce the skew to less than 1 ps.

By modifying the simulations to eliminate capacitance mismatch from the MTJs to the drain and top gate terminals it is possible to get a clearer understanding of the relationship between the other mismatches. In Fig. 4.17 the pairs plot has been reduced to the mismatches to

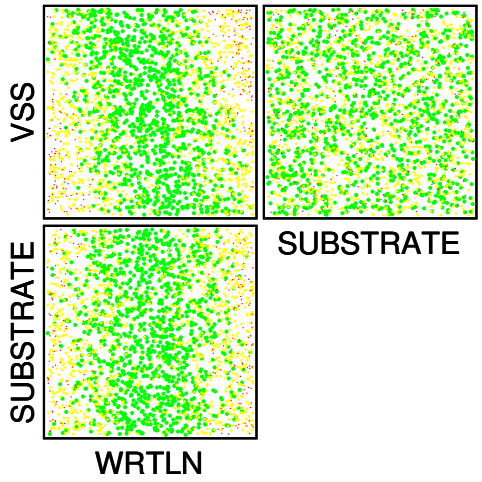

Figure 4.17: Pairs plot, capacitance mismatch and delay skew

WRTLN, VSS, and the substrate. Now the green dots represent delay skews less than 2 ps and the red dots are skews greater than 5 ps. In this case it appears that the mismatch to WRTLN is most highly correlated to delay skew, with a small mutual correlation with VSS. In these simulations WRTLN is driven to a low voltage with large NMOS transistors, so it is essentially equivalent to the VSS node. However, the parasitic capacitance to WRTLN is roughly four times greater than to VSS, so large changes in the mismatch to VSS will only offset small changes in the mismatch to WRTLN.

Optimizing parasitic capacitances at the MTJ nodes is obviously critical, but it is also important to examine the effect of parasitics at the other nodes involved in resolving the state of the MTJs, namely the nodes of the DICE latch, Y0 through Y3. Only one example of this optimization process will be presented here, regarding the capacitance from Y0 through Y3 to VSS. Note that the random variables in the Monte Carlo simulation are now the actual

capacitance values rather than a capacitance mismatch value. If 2500 simulations are performed, allowing the four relevant capacitance values to vary ± 30% from their nominal value of 0.5 fF, then the pairs plot reveals the correlations shown in Fig. 4.18. More correctly,

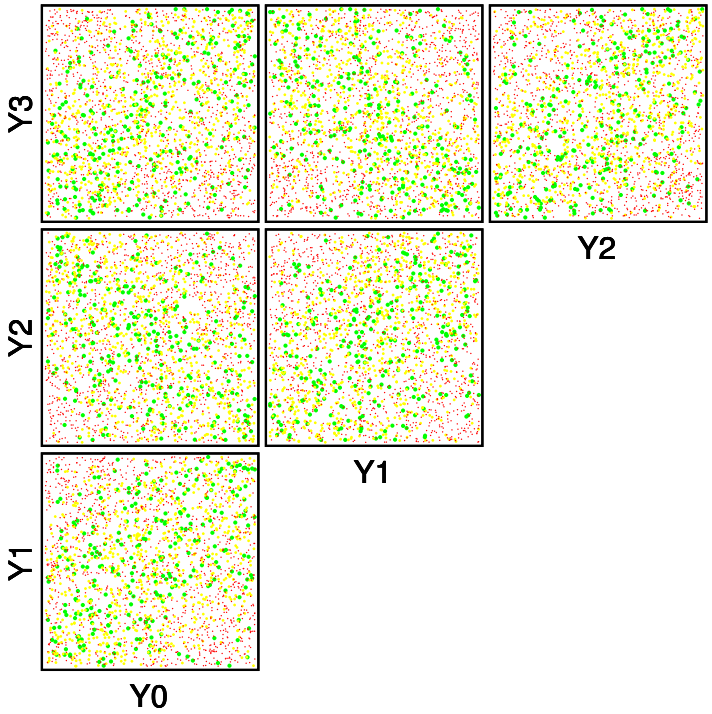

Figure 4.18: Pairs plot, capacitance from Y0-Y3 to VSS and delay skew

the pairs plot reveals the *lack* of correlations. The eye can discern faint tendencies for the green dots to be aligned with the diagonals of the individual plots, but the correlations are weak at best.

Better insight can be gained by recalling that nodes Y0 and Y2 are a redundant pair, as are nodes Y1 and Y3. Using algebraic combinations of the nodes within a redundant pair, or across the redundant pairs, yields the pairs plot shown in Fig. 4.19. The top row of plots in Fig. 4.19 shows correlations between the algebraic combinations of nodes within the redundant pairs. A very strong positive correlation is seen between the sum of capacitances from Y1 and Y3 to VSS and the sum of capacitances from Y0 and Y2 to VSS. Proceeding across the top row of plots, we see that there is little or no significance in the difference

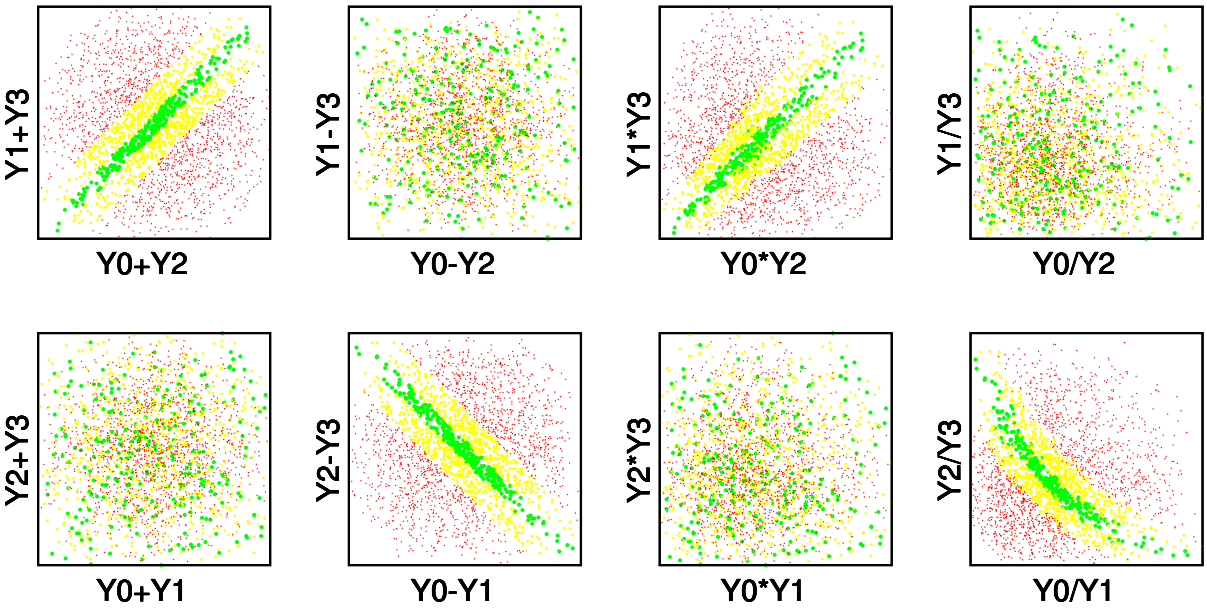

Figure 4.19: Pairs plot, paired capacitance from Y0-Y3 to VSS and delay skew

between capacitances for nodes within a redundant pair, but another strong positive correlation to their product. The rightmost plot evaluates the correlation between the ratios of capacitance values within nodes pairs, and little correlation is apparent.

The bottom row of plots in Fig. 4.19 looks at correlations between algebraic combinations of the capacitances for nodes that are not in the same redundant pair. These plots illustrate complementary relationships to those seen in the top row. Now there is little significance to the sums of the capacitances and a strong *negative* correlation to their differences. This is consistent with the upper left plot, since $Y1+Y3 = Y0+Y2$ implies $Y2-Y3 = -(Y0-Y1)$. Similarly, the strong correlation shown for $Y1 \times Y3 = Y0 \times Y2$ suggests that there should also be a strong correlation for $Y2/Y3 = 1/(Y0/Y1)$, as shown in the rightmost plot on the bottom row. The message to the circuit designer is that matching the sums of the capacitances for the redundant node pairs is most important, and that some difference between the capacitances to nodes within a redundant pair is tolerable.

Having optimized the circuit based on the well-understood parasitic elements it is important to verify that the design is robust, in the sense that it will continue to function properly over a

wide range of MTJ characteristics. The behavior of the shadow latch (with Q5C through Q8C added) as a function of the nominal MTJ resistance and capacitance is shown in Fig. 4.20. The TMRR was 30% for these simulations, and the data shown here is for a latch resolving to the '0' state. The ΔV that develops between BG02 and BG13 is shown on the left and the

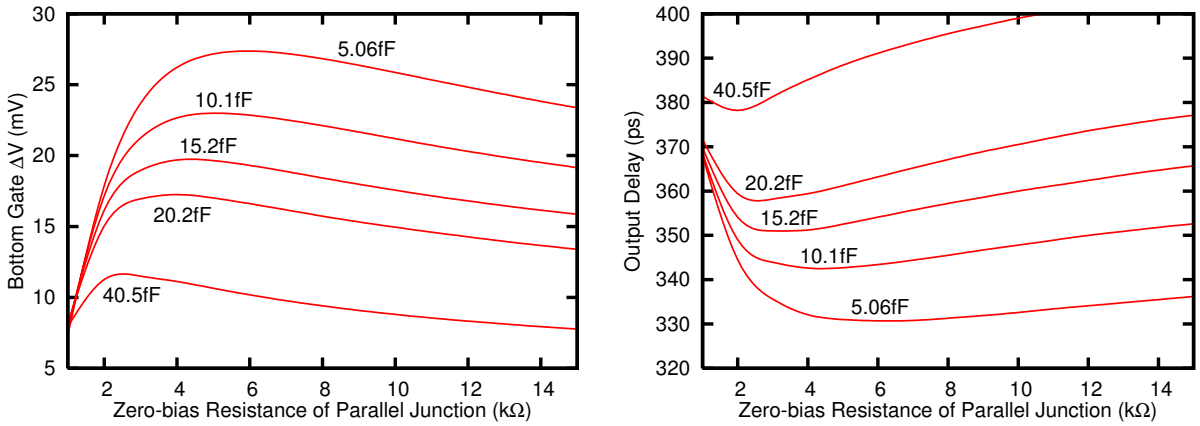

Figure 4.20: Shadow latch characterization, third layout

output transition delay is shown on the right. These curves should be compared to Fig. 4.11 for the circuit without added capacitance. Note that the peak ΔV values have been increased by at least a factor of three, and the added capacitance has not significantly degraded the output delay.

The latch can also be characterized as a function of the TMRR, as shown in Fig. 4.21. Results are given for two values of the nominal MTJ resistance, 5 kΩ in the bottom two graphs and 10 kΩ in the top two graphs, and for the same MTJ capacitance values used previously. The circled points are those combinations of TMRR and MTJ capacitance where the latch is functional, and it is not surprising that failures occur for low TMRR values. Note that in this case the results given are for the latch resolving to the '1' state, which are somewhat poorer than the results for the '0' state due to the residual imbalance in the cell. Obviously, higher TMRR values will improve the operation of the latch but it performs well for all values over 30% and can be used with a TMRR below 20% in some cases.

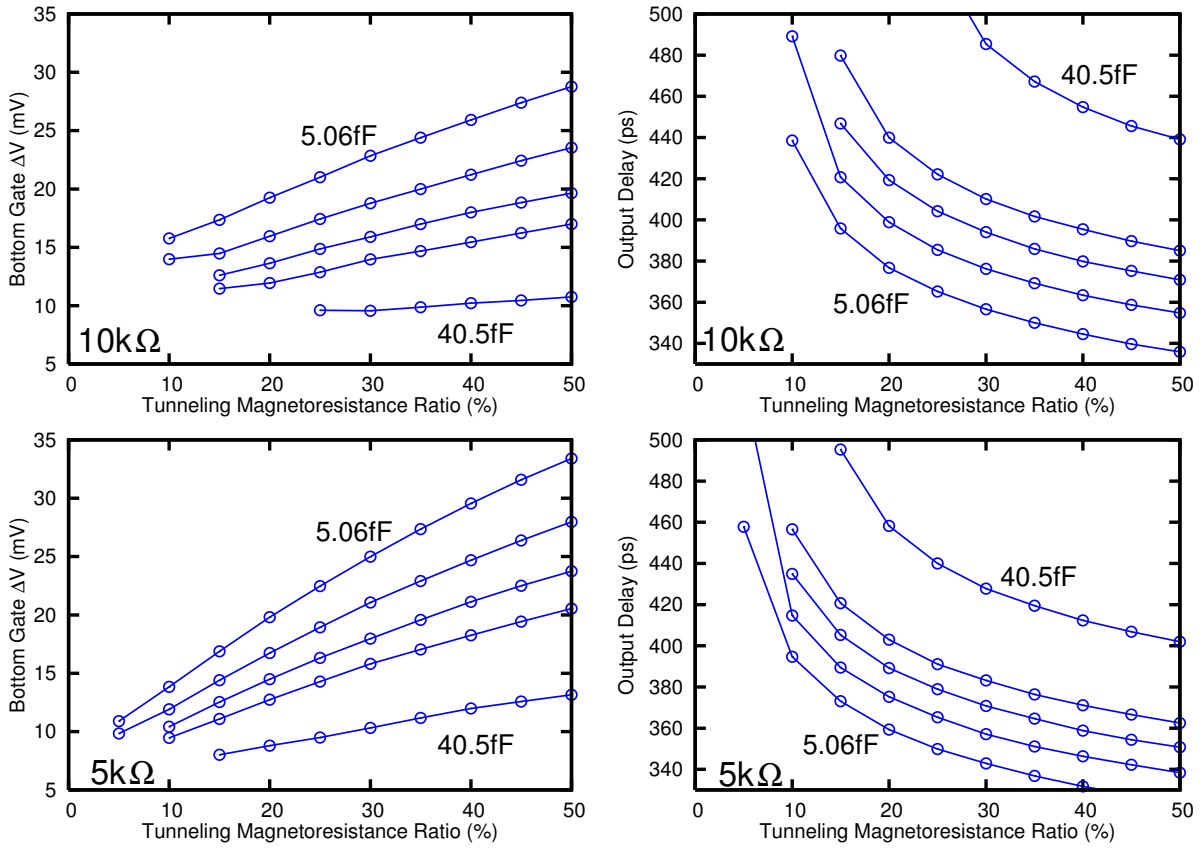

Figure 4.21: Shadow latch TMRR characterization, third layout

4.4.2 Read Operation With Active WRTLN

All of the simulations and optimizations described above assumed that the common node of the two MTJs was essentially fixed at ground. Of course, this node is actually connected to WRTLN and if the logic state stored in the MTJs must be retrieved immediately after they are programmed then WRTLN will not be grounded. In this case WRTLN is likely to be around $V_{DD}/2$, and the bottom gate will be the sink of current flowing through the MTJ. The bottom gate forms a PN junction to the source, so at some point this junction will be forward biased and will limit the voltage rise. This behavior can be seen at the left of the top waveform plot of Fig. 4.22, where WRTLN is about 1.6 V and the bottom gate nodes, BG02 and BG13, are at roughly 0.9 V. The voltage across the MTJs is therefore approximately 0.7 V. With this bias across them there is very little difference between the resistance of the parallel and antiparallel

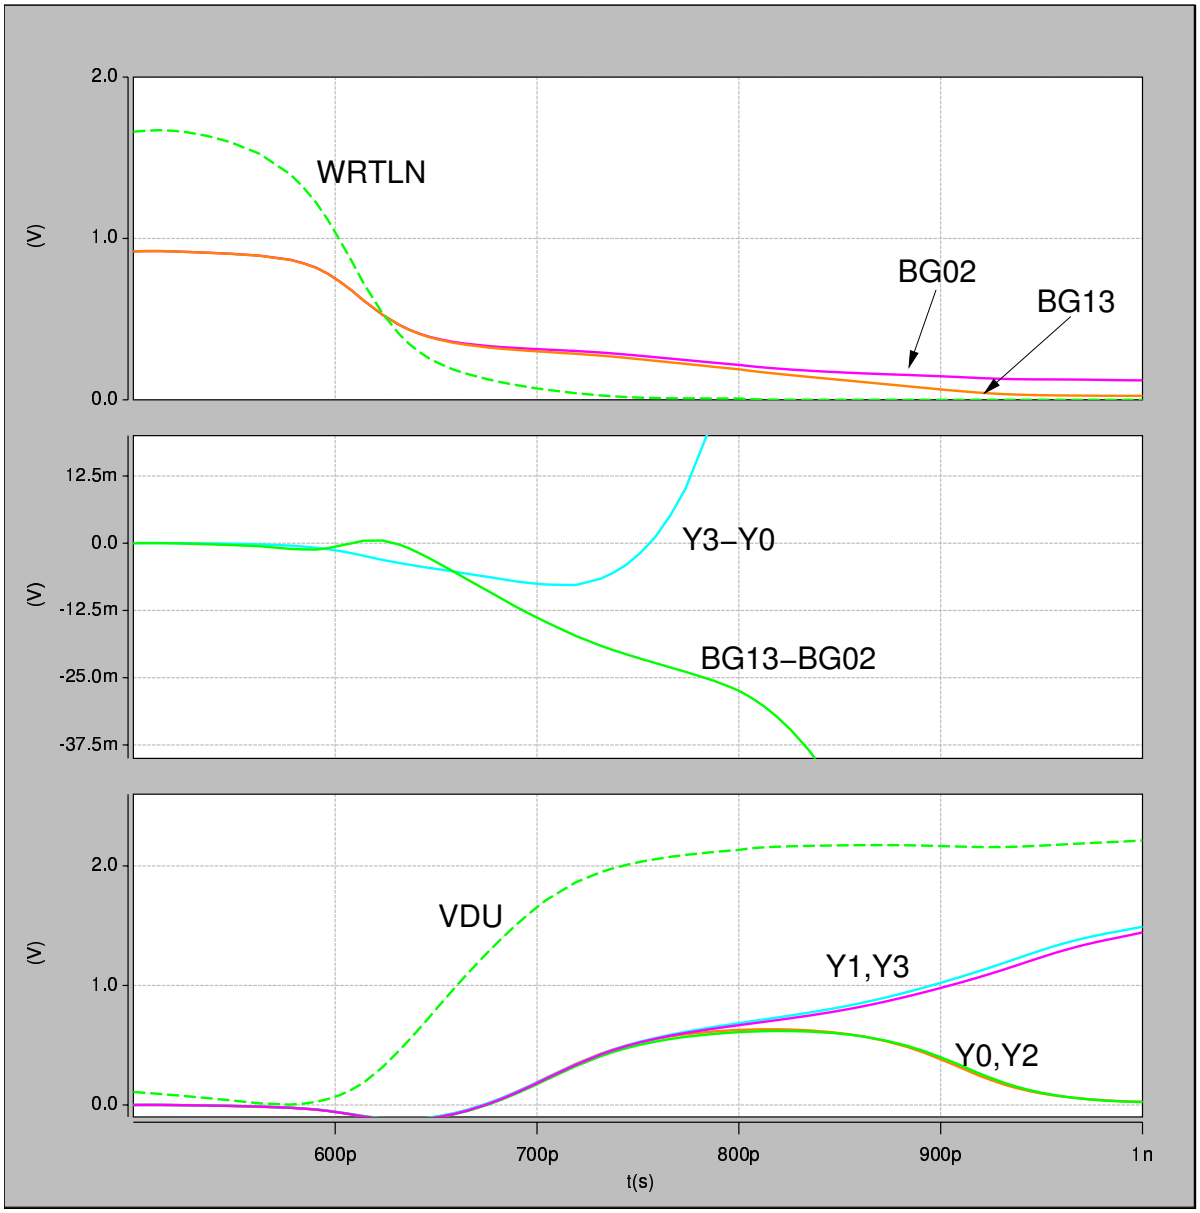

Figure 4.22: Shadow latch simulation, program/read scenario

junctions, so the bottom gate nodes rise to essentially the same voltage. This figure should be compared to Fig. 4.10 and Fig. 4.13, where the bottom gate nodes were initialized to ground.

Unfortunately, the forward-bias characteristics of the bottom gate have not been thoroughly characterized for the FlexFET since the transistors are not typically used in this mode. For the simulations discussed here a diode model taken from a P+/N-Well diode in a commercial bulk-CMOS process has been used to model the forward bias behavior. Therefore, these

simulation results must be considered only an approximation of the actual circuit behavior until a more complete model of the bottom gate junction is available. No attempt will be made to perform a detailed optimization based on these incomplete models.

Nevertheless, the simulation results in Fig. 4.22 reveal the important differences in this mode of operation. After the MTJ programming step is completed WRTLN is again brought low and VDU rises to begin the step of reading the MTJ states. Note that when VDU rises the bottom gate nodes are already biased at several hundred millivolts, a much higher level than would result from the capacitive coupling to the front gates. The ΔV that develops between the bottom gates is primarily due to the discharge of these nodes through the MTJs, and is at least as large as what was achieved solely by capacitive coupling.

The critical conditions for the success of the shadow latch are the same, whether the bottom gates are precharged or grounded at the beginning of a read operation. The bottom gates must be driven to the same potential before the read operation, and a significant ΔV must be allowed to develop before the circuit reaches a metastable state. The bias dependence of the TMRR dictates that the difference between the bottom gate voltages and WRTLN should not exceed a few hundred millivolts while attempting to resolve the state of the MTJs, and the danger of dielectric breakdown suggests that the voltage across the MTJs should not exceed approximately one volt at any time. To take advantage of the highest small signal gain in the FlexFET it is important to maintain the bottom gate voltages below about 250 mV when attempting to resolve the MTJ states.

4.5 Bulk CMOS Implementation

There is an opportunity to extend this work by applying the same circuit concepts to a bulk CMOS process, where the MTJs are connected to the lightly doped well that is the transistor body rather than to a bottom gate. A notional schematic for a shadow latch in a

P-substrate/N-well CMOS technology is shown in Fig. 4.23, which can be compared to the FlexFET circuit in Fig. 4.2. The bulk CMOS circuit is essentially the complement of the SOI

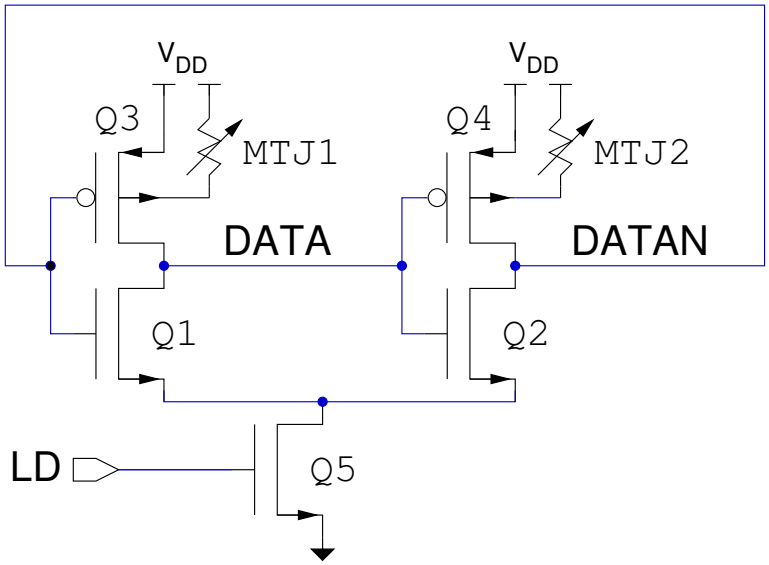

Figure 4.23: Bulk CMOS shadow RAM cell

circuit since it is now the PMOS transistor thresholds that are controlled by the MTJs. Note that the signal that controls the precharge step, LD, has also been inverted and is now asserted high.

Unfortunately, using bulk CMOS rather than SOI is likely to result in a significantly larger layout. Unlike typical N-wells, the two N-wells that are connected to the two MTJs are electrically isolated from the supply voltage and are therefore categorized as “hot” wells. Hot wells increase the risk of latchup so the required minimum spacing between these wells and any other diffusions is typically much greater than the spacing required for wells tied to V_{DD}. In the extreme case it would be necessary to add guardrings around the hot wells.

A further disadvantage of a bulk CMOS implementation is that these processes are typically engineered to reduce the body effect, which is the degree to which the body-to-substrate voltage modifies the transistor threshold voltage. Consequently the small signal gain from the body terminal to the drain is smaller, as shown in Fig. 4.24. These simulations were performed with models for the Taiwan Semiconductor Manufacturing Company (TSMC)

0.18-μm process with a supply voltage of 1.8 V. The PMOS transistor was made much larger than the NMOS transistors to overcome its inherently lower transconductance and provide increased gain. In these simulations, the W/L ratio (in μm) is 6.0/0.24 for Q1, 1.0/0.24 for Q2, and 1.0/0.18 for Q3. The maximum gain achieved for this circuit as only about 5 V/V, compared to nearly 14 V/V for the FlexFET circuit in Fig. 4.6.

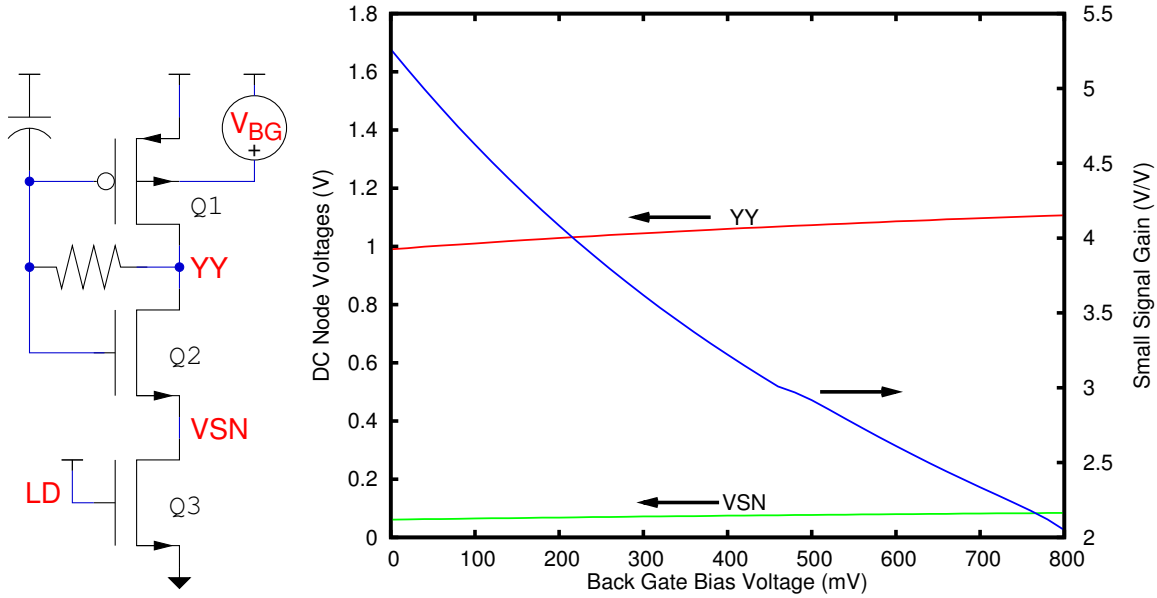

Figure 4.24: Bulk CMOS back gate small signal gain

Chapter 5
One-wire Programming

5.1 Overview

The circuitry and structures used to program an embedded shadow latch are much different than those used in conventional bulk magnetic RAMs. Bulk magnetic RAM uses two orthogonal write lines, allowing conventional X-Y addressing of individual bits [22, 176]. The net magnetic field at the intersection of a row write line and a column write line is the vector sum of the individual fields, which must be large enough to reorient the free layer. Although this arrangement allows for dense packing of the bits, it also means that a portion of the write field is applied to all of the bits on the selected row and all of the bits on the selected column. Therefore, the design of these RAMs requires a careful compromise between minimizing the necessary write field (to reduce energy consumption) and preventing write disturbances to the half-selected bits.

For an embedded MRAM, where the storage bits are intermixed with logic gates, the X-Y addressing scheme must be abandoned. Instead, a "one-wire" programming technique has been developed, as shown in the notional diagram of Fig. 5.1 [174]. (This may be recognized as the "H bridge" circuit used to reverse the rotation of d.c. motors, and the concept is also used in high-power bridge amplifiers.) This figure presents a top view of the MTJ elements placed on top of the metal write line. The idealized switches illustrate how steering the current through the write line accomplishes the programming of a '1' or '0' state. As shown on the left side of the figure, current flows first from left to right under `MTJ1` and then from right to left under `MTJ2`. The green lines indicate the direction of the resulting magnetic field.

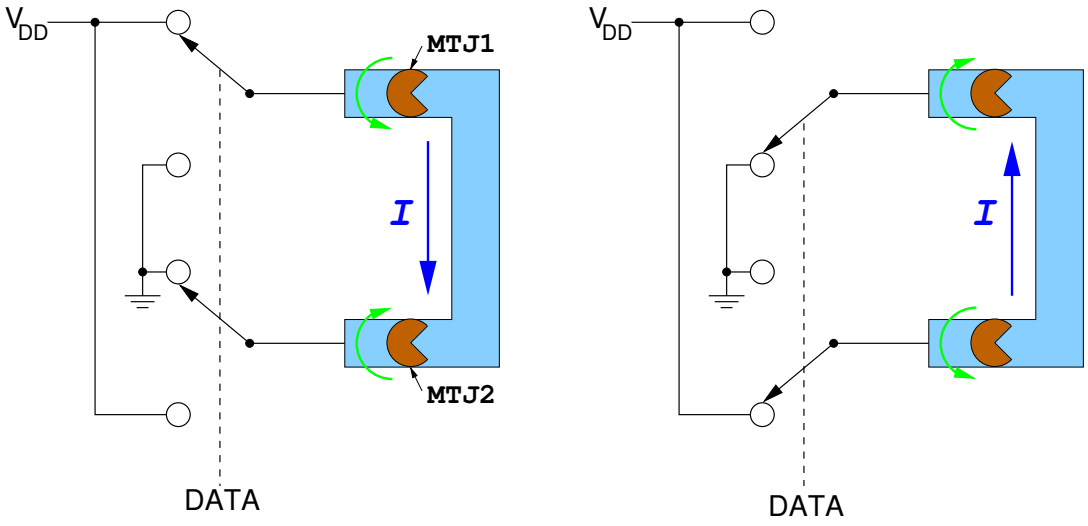

Figure 5.1: One-wire write scheme

Since the orientation of the pinned layer is fixed during manufacturing and is the same for all MTJs, we must change the physical direction of current flow along the write line to program the two MTJs to opposite states, as shown on the right side of Fig. 5.1. Note that the figure illustrates the MTJs as being on top of the word line for clarity but constructing the write line directly on top of the MTJs is essentially equivalent, as long as the sequence of layers in the MTJ is such that the free layer is nearest the write line.

A variation of the approach shown in Fig. 5.1 was previously proposed, but it utilized two word lines where one of them was folded to allow programming of a differential pair [158]. The one-wire programming described here takes advantage of the close physical proximity between the MTJ free layer and the write line, which greatly increases the efficiency of the current flowing in that line but similarly diminishes the ability of a second word line to contribute to the programming field.

In bulk MRAM cells the hard and easy axes of the free layer are typically aligned with the two orthogonal write lines. Consequently, the direction of the effective magnetic field is turned approximately 45° away from either write line. This configuration produces faster and

more reliable switching of the free layer, with better selectivity, then if the magnetic field were directly aligned with the easy axis [56]. In the one-wire programming scheme the net magnetic field is always at right angles to the write line, so the physical orientation of the free layer's easy axis is rotated 45° to recreate the desired angle [177].

Suppose that the situation illustrated on the left side of Fig. 5.1 causes `MTJ1` to have a parallel magnetic state while `MTJ2` is forced to the antiparallel state, which we can arbitrarily declare to be the '0' state of the memory bit. The switches shown in the diagram are implemented with transistors, and can reverse the overall current flow through the write line to program the opposite logical state of the memory bit. Reversing the current flow will reverse the direction of the magnetic fields as shown on the right side of Fig. 5.1, so that `MTJ1` is now in the parallel state and `MTJ2` is antiparallel, which would represent the '1' logical state of the memory bit. For a shadow latch the DATA signal that controls the switch is connected to the DATA signal in Fig. 4.2 so the state of the static RAM portion of the cell will be stored in the MTJs, allowing a "snapshot" of the current information in the integrated circuit to be saved for quick recovery after power is interrupted. The same concept can be used to quickly change the function of programmable logic if the DATA signal that controls the write line is distinct from the state of the static RAM so that a different personalization of the integrated circuit can be loaded in to the MTJs without disturbing the static RAM.

An important distinction between a true shadow RAM application and other applications should be emphasized here. In a shadow RAM the node voltages in the latch circuit cannot be disturbed until after the MTJs are written, because these voltages define the new data to be stored. Therefore, the signal that precharges the latch cell must be distinct from the signal that activates the write line drivers, so that the precharge does not begin until after the new data is written. For other applications, where the data to be stored is independent of the current state of the latch, these two signals can be combined so that the precharge and the write operation are initiated simultaneously.

The MTJs are actually quite thin (tens of nanometers) and placed directly on top of the write line, so that the magnetic field in their vicinity is almost entirely parallel to the top surface of the write line. This achieves the most efficient conversion of write current to useful magnetic field (see Section 1.3). The distance from the MTJ's free layer to any conductor other than the write line is much greater, so it is not difficult to prevent fields produced by adjacent wires from accidentally programming the cells.

Although Fig. 5.1 shows the write line connected directly between V_{DD} and ground, in practice a number of write line circuits can be connected in series and share the same programming current. Current would enter the write line circuit (from the previous bit in the serial string) at the node labeled V_{CC} and would exit the write line circuit (to the next bit in the serial string) at the node shown connected to ground in the figure. The limit to the number of bits that can be strung together in this fashion is determined by the voltage drop across each write line circuit, which is in turn determined by the size of the transistors that steer the current and the resistance of the write line itself. This approach will be explored in detail in Section 5.3.

The write operation, where the MTJ free layers are reoriented to store information, largely determines the energy efficiency of magnetic memories. Relatively large currents must pass through the write line and its associated switching transistors to create the necessary magnetic field. Micromagnetic simulations predict that a current of approximately 20 mA is needed for a typical aluminum or copper line [177], but this can be reduced significantly by cladding the write line with a ferromagnetic material [178, 179]. Fortunately, the write pulse is quite short compared to other nonvolatile memory technologies, lasting less than 5 ns.

5.2 Single-Bit Write Line

The simplest implementation of the write line circuitry for a single latch is shown in Fig. 5.2, with the state table shown in Table 5.1. The two NAND gates, G1 and G2, will have high

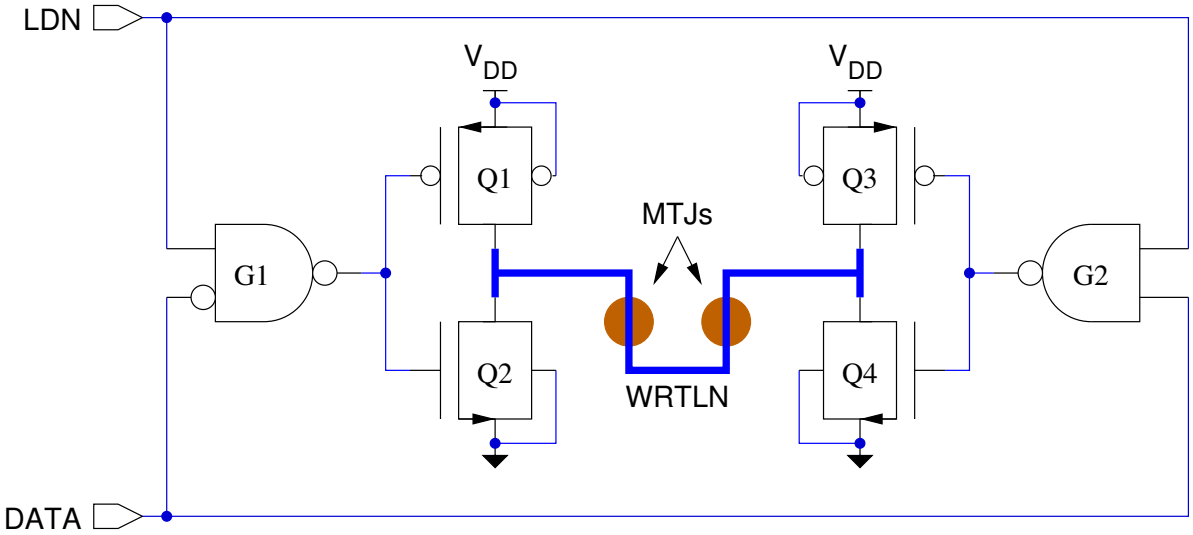

Figure 5.2: One-wire programming schematic

outputs whenever LDN is low to insure that neither of the PMOS transistors, Q1 or Q2, will be enabled but both of the NMOS transistors are enabled. When LDN is deasserted (high) the direction of current flow through the write line is controlled by the DATA signal. When DATA is high the output of G2 will fall while the output of G1 remains high, enabling transistors Q3 and Q2 to drive current from right to left through WRTLN. On the other hand, if DATA is low then Q1 and Q4 are enabled to drive current from left to right.

Operation	LDN	DATA	Left End	Right End	Condition
Read	0	0	0	0	WRTLN grounded
	0	1	0	0	
Write	1	0	1	0	Current flow left-to-right
	1	1	0	1	Current flow right-to-left

Table 5.1: Typical write line operation

One embodiment of a one-wire programming circuit is shown in Fig. 5.3, using FlexFET transistors. This layout plot shows only the wiring on Metal1 and Metal2, plus the vias that

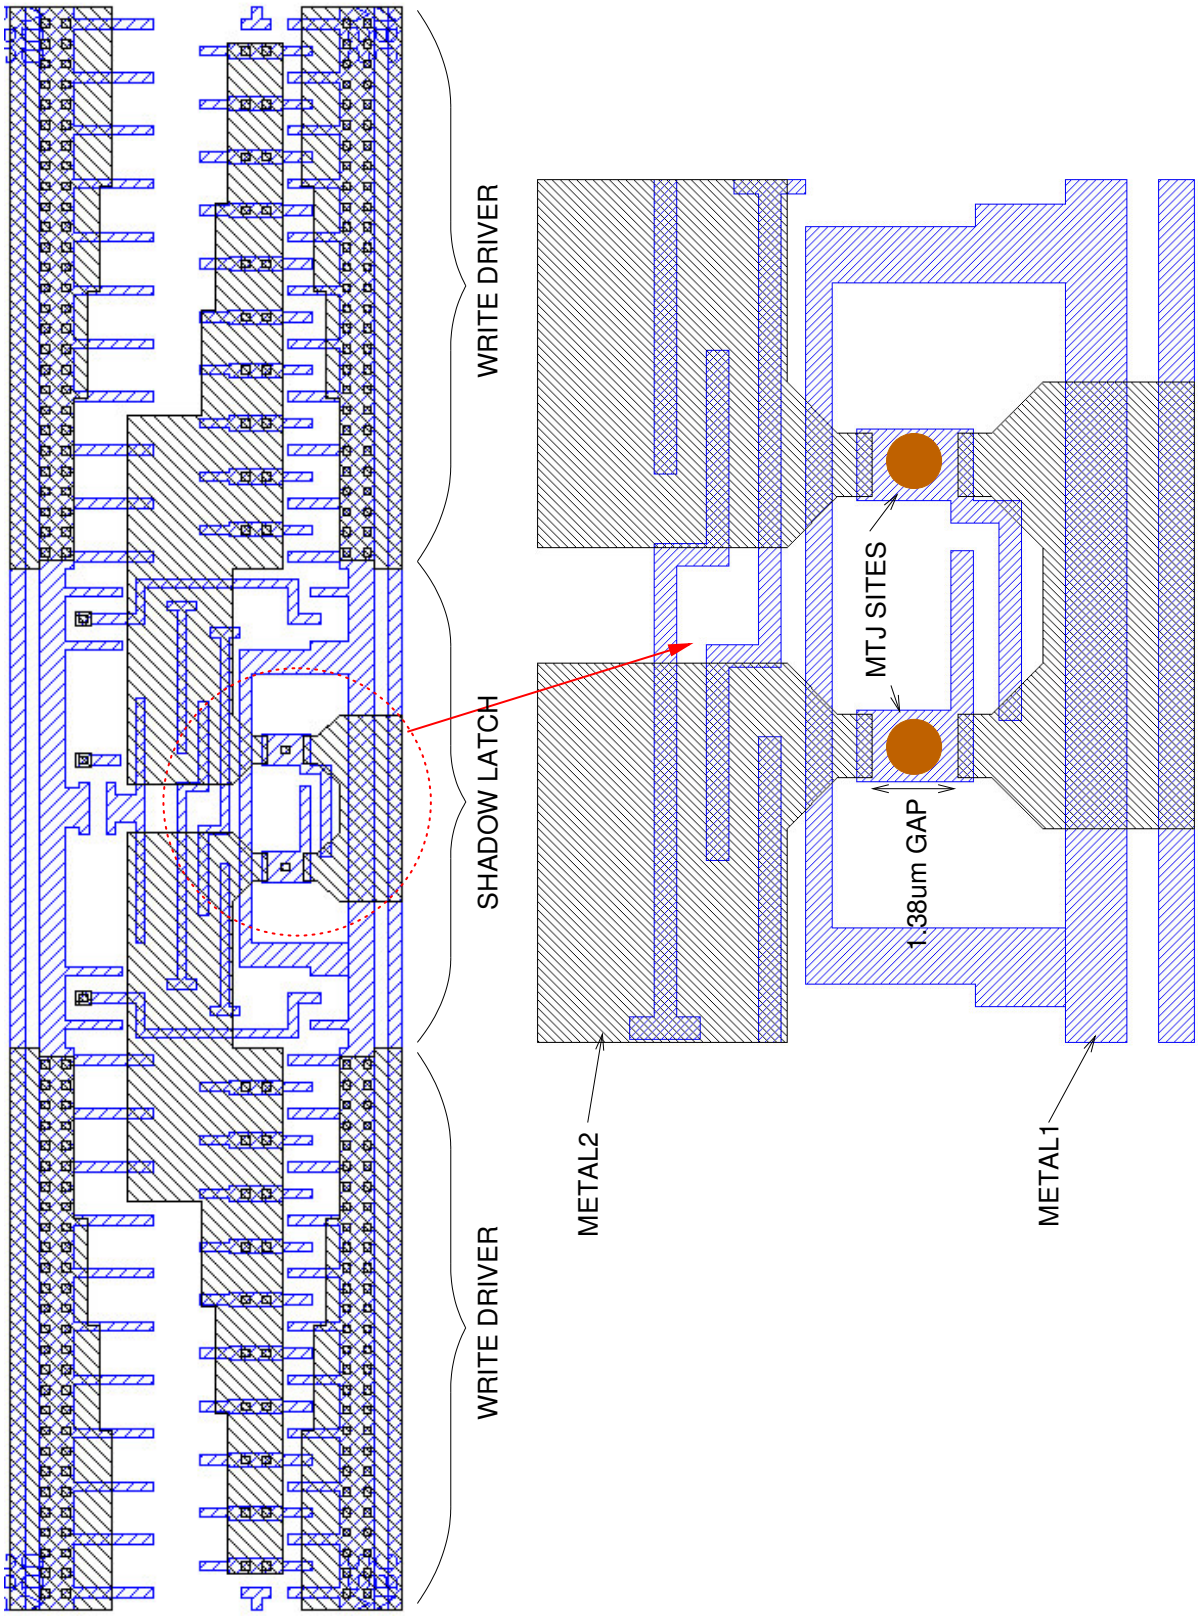

Figure 5.3: One-wire programming layout

connect these two layers. About two-thirds of the layout comprises the large transistors (`Q1` through `Q4`) that must source or sink the write line current, split between the right and left sides, and this example also omits gates `G1` and `G2`. The circuitry of the latch itself is located between the drivers. Each of the write drivers is electrically equivalent to a single PMOS transistor and a single NMOS transistor, but each of these transistors is divided into twenty physical transistor "fingers" with a distinct gate, connected in parallel. The transistor drain islands are shared between a pair of fingers so there are 10 distinct connections from the transistor drains to the write line structure on Metal2, each capable of providing a minimum of 2 mA through two vias.

The calculated current density through the write line structure on Metal2 is shown in Fig. 5.4 [8]. The layout is symmetrical so only one half of the layout is shown in this figure, and the model assumes that Metal2 is connected across the gap where the MTJ will be placed. In this model the write line is specified to have a thickness of 0.65 μm and a width of 1.02 μm. The total current injected into the structure is 20.0 mA, for a peak current density of 3.01 MA/cm^2 (30.1 GA/m^2). As expected, the current density is below 1 MA/cm^2 over most of the larger metal features.

Figure 5.4: Write-line current density, MA/cm^2

The corresponding voltage drop along the write line is illustrated in Fig. 5.5. The structure is assumed to be aluminum with a resistivity of 30 $n\Omega \cdot m$ (46 $m\Omega/\square$). Thus, the voltage drop is

roughly 1 mV/□. Clearly, most of the voltage drop (and power consumption) occurs in the driver transistors themselves and not in the metal lines.

Figure 5.5: Write-line voltage drop, mV

Simulations of the circuit netlist with all extracted resistors and capacitors were performed at the three corner cases shown in Table 5.2. The HSPICE simulations confirm that there is very little voltage drop along the write line wiring. For the nominal case the highest interconnect voltage (at the drains of the driving PMOS transistors) is 1.683 V and the lowest voltage (at the drains of the driving NMOS transistors) is 1.641 V, for a total drop of 42 mV and a net resistance of 1.14 Ω. Checking the voltage drop across the parasitic resistances reveals that almost all of it occurs in the contacts between metal 1 (MM1) and metal 0 (MET0). Based on these results, redundant contacts should be added to reduce the network resistance and the current density in those contacts.

Conditions	V_{DD}	Temp	FlexFET Models	Current	WRTLN voltage
Worst Case	2.25 V	125 °C	Slow	25.7 mA	1.29 V
Nominal	2.50 V	25 °C	Typical	36.7 mA	1.66 V
Best Case	2.75 V	−55 °C	Fast	48.1 mA	1.99 V

Table 5.2: Write line driver corner cases

The circuit is designed to provide the minimum necessary write current, with some additional margin, under the “worst case” conditions of low supply voltage, high temperature, and slow (weak) transistors. Note that the available write current is about 25% higher than the expected

requirement, so it may be possible to reduce the size of the driver transistors by 20% once the actual current requirement is well understood. The nominal voltage on WRTLN under these conditions is about 15% higher than $V_{DD}/2$ so the ratio of PMOS:NMOS transistor width is also somewhat high. For purposes of specifying the maximum power consumption and determining the highest current density, the "best case" conditions are used. When the supply voltage is increased to 2.75 V, the temperature is lowered to -55 °C, and the fast (strong) transistors models are used the WRTLN current nearly doubles compared to the worst case simulation. Note that the power consumption for these conditions is 132 mW during the programming of a single bit.

5.3 Serial Write Lines

Clearly, the power consumption of the write line will limit the number of MTJs that can be programmed simultaneously. In certain applications, such as the personalization of programmable logic, it may be perfectly acceptable to program no more than 8 or 16 latches at a time. In this case the write line driver circuit is modified to allow the same current to program multiple latches by serially connecting several write line drivers. This requires a significant reduction in the saturated V_{DS} of the driver transistors, suggesting that the widths of the driver transistors would need to be increased proportionally. However, if there are multiple drivers in series then it is also true that they will experience different source voltage levels and hence different V_{GS} voltages.

Consider the example shown in Fig. 5.6, which shows three write line circuits (from Fig. 5.2) in series so that three latches can be programmed with the same current, and for the sake of discussion only the enabled driver transistors are shown. If we assume that V_{DS} is the same for every transistor ($V_{DD}/6$) and ignore wiring resistance, then the WRTLN node in `LATCH 1` (W1) is at $\frac{5}{6}V_{DD}$, at $\frac{1}{2}V_{DD}$ in `LATCH 2` (W2), and at $\frac{1}{6}V_{DD}$ in `LATCH 3` (W3). The connection

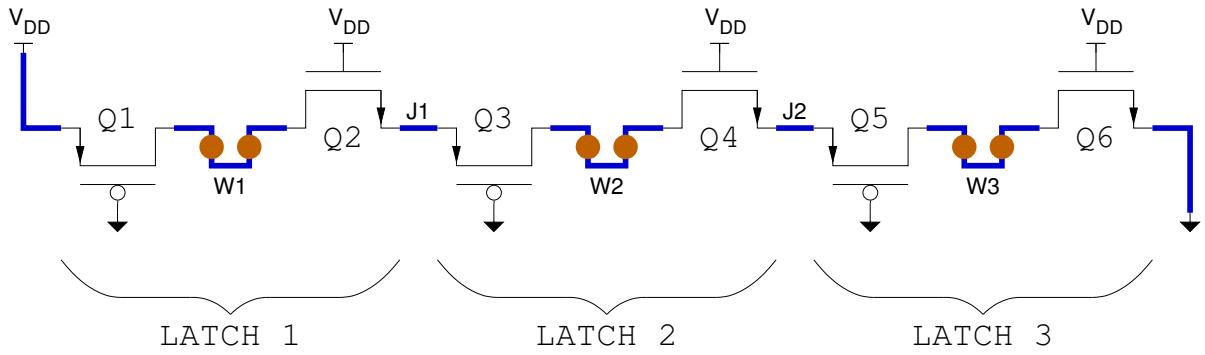

Figure 5.6: Series connected write line drivers

between LATCH 1 and LATCH 2, J1, is at $\frac{2}{3}V_{DD}$ and the connection between LATCH 2 and LATCH 3, J2, is at $\frac{1}{3}V_{DD}$. Consequently, V_{GS} for Q2 and Q5 is only $\frac{1}{3}V_{DD}$ and these transistors must be very large to provide a small V_{DS}. (Note that all transistors are assumed be FlexFET devices with the bottom gate tied to the source.) If Q2 is replaced with a PMOS transistor with its gate at ground instead of V_{DD}, then this new PMOS transistor will have a V_{GS} of $\frac{5}{6}V_{DD}$ and it should be possible to use a smaller transistor to achieve the same V_{DS}. The trade-off for replacing Q5 with an NMOS transistor is more clear cut, and the higher mobility for electrons will allow a much smaller NMOS transistor to be used. Simulations in HSPICE were used to determine the transistor widths needed to obtain an I_{DS} of 25 mA for this case, where V_{DS} for all transistors is forced to be $V_{DD}/6$, with the results shown in Table 5.3. The intuitive analysis is confirmed, and it is clearly impractical to use a PMOS for Q5. Smaller but very significant reductions in circuit area result from replacing Q2 with a PMOS transistor and Q3 with an NMOS transistor. Note that the Table 5.3 lists the widths of the *enabled* transistors only, and for each enabled transistor there is an equally-sized disabled transistor that would be enabled instead to reverse the direction of current flow, so the total required transistor width is actually twice what is listed.

Of course, the constraint that all V_{DS} values be equal is artificial, as the magnitude of the voltage on the WRTLN has no effect on the MTJ programming operation. Without this constraint the total transistor width can be reduced even further. A genetic optimization was

used to determine the widths of six transistors in series that would allow an I_{DS} of 25 mA if the gates of all PMOS transistors were connected to ground and the gates of all NMOS transistors were connected to V_{DD}. The "genes" of the population were the transistor widths, and the fitness of any given solution was determined from two parameters, the I_{DS} of the transistor string and the sum of their widths (in μm):

$$\text{fitness} = \begin{cases} 1000000 & \text{if } I_{DS} < 25 \text{ mA}, \\ \sum_n W_n & \text{if } I_{DS} \geq 25 \text{ mA}. \end{cases} \tag{5.1}$$

The term "fitness" is somewhat misleading because the goal of the optimization software is actually to minimize the fitness function. Therefore, any solution that did not supply 25 mA was extremely unfit, and solutions that supplied at least 25 mA were more fit if they used less total transistor gate width. The genetic algorithm is not guaranteed to find a global minimum, and the process often results in several sets of dissimilar "genes" that are almost equally fit.

The results of several optimization runs for this problem are shown in Table 5.4, with transistor gate widths shown in micrometers. The first row presents a solution for the case shown in Fig. 5.6, where each latch has PMOS transistors on the high side of WRTLN and NMOS transistors on the low side. The second row shows a solution obtained by replacing `Q2` with a PMOS transistor, and the third row shows a solution obtained by also replacing `Q5` with an NMOS transistor. The fourth row gives a solution when `Q3` is then replaced with an NMOS transistor, to result in the configuration that was optimal for the fixed V_{DS} case but now with

Q1		Q2		Q3		Q4		Q5		Q6		Total
P/N	Size	P/N	Size	P/N	Size	P/N	Size	P/N	Size	P/N	Size	Width
P	233	N	646	P	392	N	113	P	2530	N	68.4	3982 μm
P	233	*P*	290	P	392	N	113	P	2530	N	68.4	3626 μm
P	233	P	290	P	392	N	113	*N*	83.8	N	68.4	1180 μm
P	233	P	290	*N*	186	N	113	N	83.8	N	68.4	974 μm

Table 5.3: Serial WRTLN driver, equal V_{DS}

approximately 14% less transistor area. However, we can go one step further. Having made Q3 an NMOS transistor it is possible to reduce the area even more by changing Q2 *back* to NMOS, resulting in the configuration shown in Fig. 5.7. Note that for the lowest area solutions the total driver transistor width can be reduced by approximately 25% (from 974 µm to 732 µm) by allowing V_{DS} to vary along the serial transistor string.

Q1		Q2		Q3		Q4		Q5		Q6		Total
P/N	Size	P/N	Size	P/N	Size	P/N	Size	P/N	Size	P/N	Size	Width
P	199	N	484	P	593	N	333	P	591	N	66	2266 µm
P	356	*P*	351	P	326	N	249	P	473	N	68	1823 µm
P	266	P	261	P	260	N	88	*N*	106	N	105	1086 µm
P	189	P	190	*N*	115	N	115	N	108	N	117	834 µm
P	123	*N*	125	N	120	N	125	N	125	N	107	732 µm

Table 5.4: Serial WRTLN driver optimization

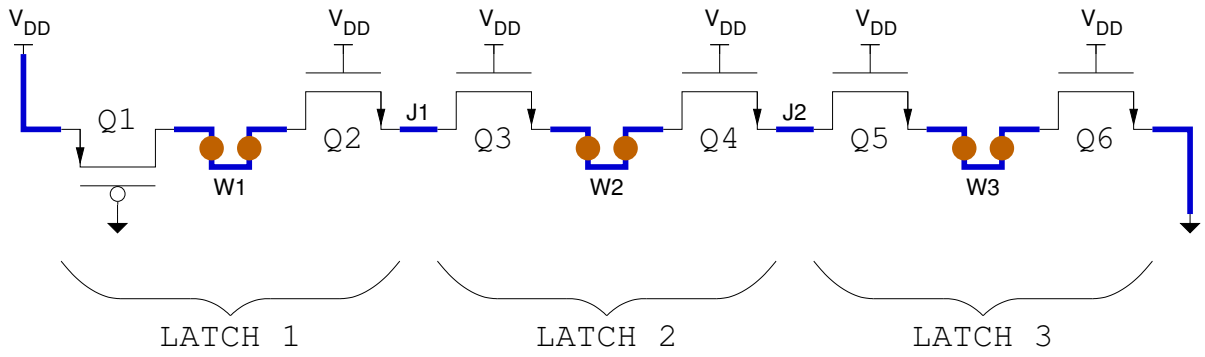

Figure 5.7: Optimized series connected write line drivers

As expected, the optimization of total transistor width results in node voltages that are significantly different than those when V_{DS} was fixed at $V_{DD}/6$, as shown in Table 5.5. Favoring NMOS transistors in the serial string has the side effect of lowering the voltages at those nodes connected to NMOS drains, while V_{DS} for the lone PMOS transistor more than doubled. As discussed earlier in regard to the shadow latch, if the voltage on WRTLN is high enough then the latch's internal nodes will be adequately discharged simply because the NMOS bottom gates are brought high. However, the WRTLN voltages for some of the latches in a serial string may not be high enough to accomplish this task and additional (small)

NMOS transistors will be needed to perform the discharge function.

Circuit	V(W1)	V(J1)	V(W2)	V(J2)	V(W3)
Equal V_{DS}	1.875	1.500	1.125	0.750	0.375
P N P N P N	1.800	1.445	1.203	1.062	0.393
P *P* P N P N	2.015	1.742	1.377	1.139	0.378
P P P N *N* N	1.926	1.517	0.850	0.466	0.220
P P *N* N N N	1.769	1.030	0.691	0.431	0.194
P *N* N N N N	1.416	0.944	0.645	0.415	0.215

Table 5.5: Serial WRTLN node voltages, three latches

Longer serial strings are possible, reducing the write energy per bit proportionally at the expense of larger transistors. Tables 5.6 and 5.7 present the results for strings of four and five latches, respectively. Unfortunately, the driver transistor area required per latch increases as the number of latches in the string increases, presenting an engineering trade-off to be evaluated for each application. For a string of three latches the total width per latch is 488 μm, for four latches it is 625 μm, and for five latches the total is 776 μm. Although this is a limited data set, it suggests that the average driver transistor width per latch, W_{LATCH}, for a series connection of N latches is $W_{LATCH} \approx N \times 150$ μm, and the total driver transistor width is $W_{TOTAL} = N \times W_{LATCH} \approx N^2 \times 150$ μm.

Circuit	Total Width
P P P N N N N N	1532 μm
P P *N* N N N N N	1318 μm
P *N* N N N N N N	1250 μm

Table 5.6: Serial WRTLN optimization, four latches

Circuit	Total Width
P P P N N N N N N N	2130 μm
P P *N* N N N N N N N	1959 μm
P *N* N N N N N N N N	1940 μm

Table 5.7: Serial WRTLN optimization, five latches

These analyses suggest that a generic schematic for an N-bit serial write line driver would resemble Fig. 5.8. Note that the control signal for the write operation, LDN, is only used in the

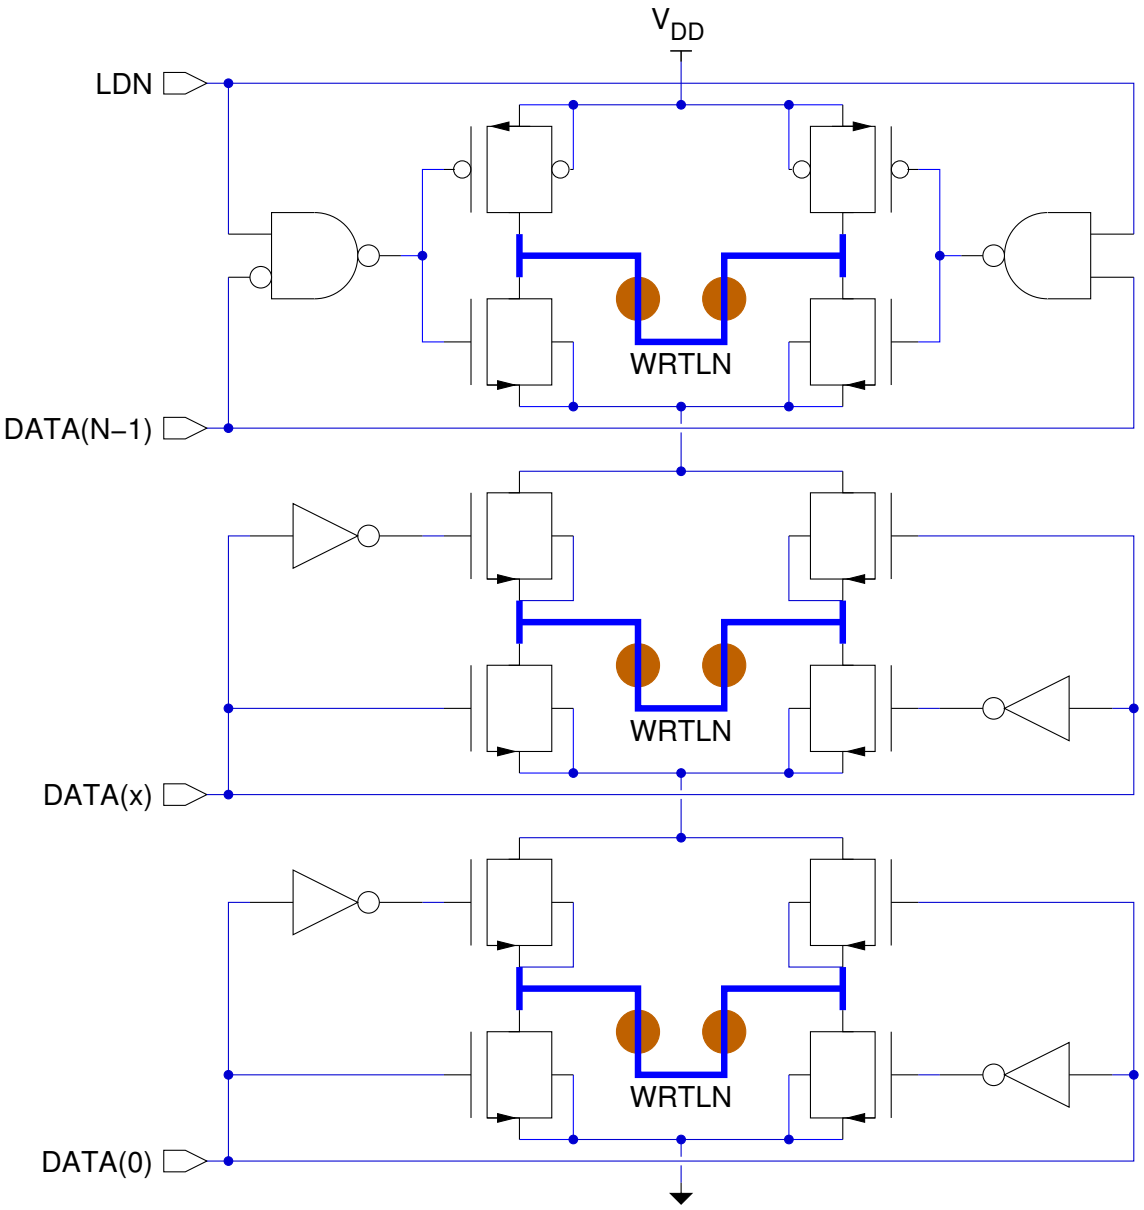

Figure 5.8: N-bit serial one-wire programming schematic

top section for bit N-1 and that this is the only section that uses PMOS transistors. When LDN is asserted low both PMOS transistors are disabled, but all of the write line sections will always have a low impedance path to ground through enabled NMOS transistors. The middle section in Fig. 5.8 is repeated N-2 times, as desired, but of course the transistor widths will vary. The bottom section is functionally identical to the middle section but has ground as its ultimate current sink.

Of course, as the width of the driver transistors increases so does the energy (per bit) needed to switch their gates and this will tend to offset the benefit of chaining the write lines. Assuming that the required write current, I_W, is a fixed design parameter then the energy consumed in writing some number N of latches can be divided into two components: the energy lost to Joule heating of the write line conductors and driver transistors, and the energy used to switch the gates of the driver transistors. The average Joule heating loss, E_J, per latch is determined by the supply voltage and the length of the write pulse so

$$E_J = \frac{V_{DD} \times I_W \times t_{WP}}{N} \tag{5.2}$$

The gate switching energy, E_G, is the energy required to charge the driver gate capacitance to initiate a write operation and then to discharge that capacitance at the end of the operation. It is determined by the supply voltage and the total gate capacitance: $E_G \approx V_{DD}^2 \times C_G$. However, the gate capacitance is proportional to the driver width, so $E_G \propto V_{DD}^2 \times N^2$. For the FlexFET technology the constant of proportionality in the previous equation is the rail-to-rail switching capacitance of 150 μm of transistor top gates, which is roughly 450 fF, so

$$E_G \approx V_{DD}^2 \times N^2 \times 450 \text{ fF} \tag{5.3}$$

This leads to the observation that when longer pulse widths are required the trade-off shifts toward preferring more latches in series in order to reduce E_J. On the other hand, if the transistor I_{DS} is low or the threshold voltage is high, requiring higher supply voltages, then E_G grows rapidly and shorter series strings are preferred. These two equations are plotted in Fig. 5.9, assuming that $V_{DD} = 2.5$ V, $I_W = 25$ mA, and $t_{WP} = 4$ ns. The configuration that consumes the least total energy per write operation has three or four latches in series, with the net usage of about 100 pJ per write operation per latch. This is roughly two orders of magnitude greater than the energy required to write a conventional CMOS flip flop at the same supply voltage, but it is at least an order of magnitude less energy than that consumed in

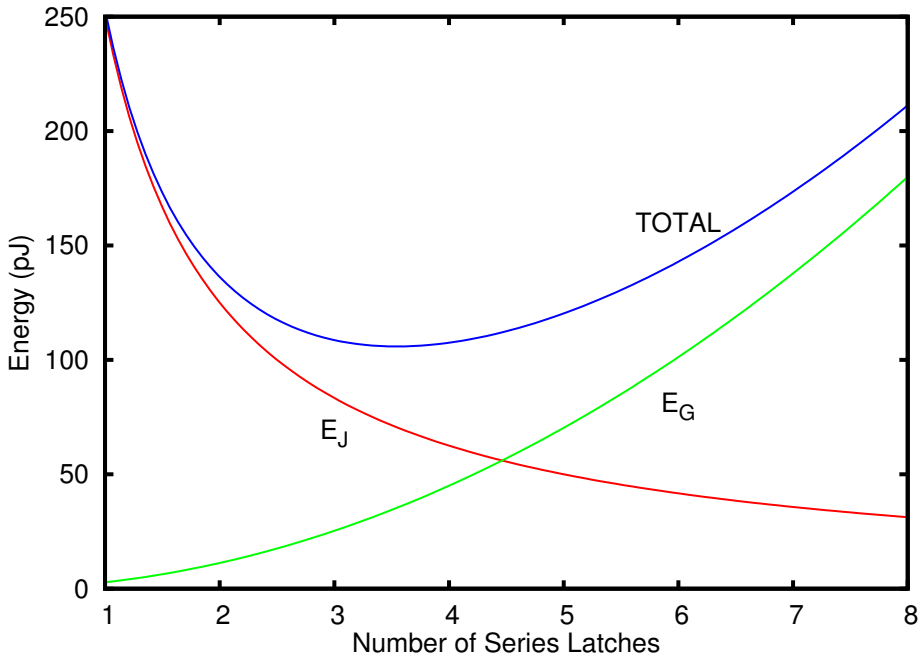

Figure 5.9: Write energy per bit for series write lines

erasing and writing a bit of flash memory [180, 181].

This reduction in write energy does not come without some cost, however. The designer must keep in mind that the area for an N-bit serial write line increases as N^2 while the area for the same number of individual write line circuits increases as N. In the example discussed above, converting three individual bits to a three-bit serial string reduces the energy consumption by a factor of about 2.5 but also increases the total write driver area by a factor of about 3. In applications such as the personalization of programmable logic, where many bits are seldom programmed, the need to reduce circuit area may tip the trade-off in favor of individual write line drivers.

The bottom gate of the FlexFET provides another avenue for reducing the size of the write line driver transistors. If the bottom gate is capacitively coupled to the top gate, then the leading edge of the pulse that enables the write driver will also provide a “boot strap” pulse to the bottom gate. The pulse on the bottom gate will lower the effective threshold voltage for the front gate and significantly decrease V_{DS} for a given value of I_{DS}. If V_{DD} is less than the

forward voltage drop of the bottom-gate-to-source junction then the two gates can be connected together, but for higher V_{DD} levels capacitive coupling should be used. The amount of coupling should be optimized to provide an effective boost of the bottom gate voltage without forward biasing it, which will require a more thorough characterization of the bottom gate junctions than is available at this writing.

Chapter 6

Reconfigurable Hardware using MTJs

As the cost of designing and manufacturing a custom integrated circuit grows rapidly, so does the desire for increasingly capable reconfigurable hardware elements. The flagship of these products is the Field Programmable Gate Array (FPGA), which devotes roughly 90% of its resources to the programming and interconnection of the remaining functional circuits. Any non-volatile memory technology must also be considered as a candidate for these applications, and magnetic tunneling junctions are no exception.

6.1 MTJs As Programmable Switches

In most cases the configuration of an FPGA can be reduced to the act of setting a very large number of switches to the "on" or "off" state. These switches may directly configure the signal interconnections on the chip by passing dynamic logic signals, or they may produce static logic values that change the function of combinational logic structures. In either case there are certain desirable qualities for the switches, as outlined in Table 6.1. The antifuse

Characteristic	Ideal	MTJ	MOSFET	Antifuse
On:Off conductance ratio	∞	< 5	10^6–10^9	$> 10^6$
Operating voltage	∞	≈ 500 mV	< 10 V	< 5 V

Table 6.1: Characteristics of programmable switches

used in programmable logic is typically formed with a very thin insulating layer between two electrodes [182]. Applying a high voltage across the insulator causes it to break down and forms a conducting bridge between the electrodes, with a resistance of hundreds of ohms.

However, the antifuse can be "erased" back to the non-conductive state by a short pulse of sufficient current, so the voltage applied to these structures during normal operation must be limited.

Clearly, the MTJ cannot be used directly as a switch. The low operating voltage means that normal digital logic signals cannot be applied to an MTJ. More importantly, the low ratio of conductance between its programmed states means that either the "open" state will have an unacceptably high conductance or the "closed" state will have an unacceptably low conductance. Therefore, additional circuitry is required for applications to programmable hardware, and proposals for using MTJs in this manner tend to fall into three categories.

6.2 MTJs Applied To Programmable Logic

The first category consists of approaches that employ exotic magnetic structures or complex write line arrangements to perform boolean logic within a single MTJ. One such proposal calls for three parallel write lines, stacked on top of each other so that all three pass over the top of the MTJ in close proximity [183]. Two of these lines are the inputs to the boolean function while the third is used to program the structure for either an AND or OR operation. Applying current to all three lines simultaneously will reorient the pinned layer to select the desired logic function. Prior to each boolean evaluation the appropriate currents must be passed through the two input lines to reset the "gate" to the antiparallel (logic 0) configuration. During the evaluation, if the currents passing in the two input lines are sufficient to flip the free layer then the gate enters the parallel (logic 1) configuration, which must be sensed as a change in resistance. This MTJ gate thus requires several metal lines in very close proximity (difficult to manufacture), carefully controlled current switching (complex write circuitry), and a resistance measurement to determine the logic state. A simplified version of this MTJ gate eliminates the third write line and the need flip the pinned

layer, but sacrifices the ability to perform OR/NOR operations [184].

The second group of proposals attempt to use the MTJ itself to create boolean operators with series and parallel combinations of MTJs that are programmed independently, with each MTJ acting as an input to the boolean expression [159, 185, 186].These schemes typically require a reference chain of MTJs to overcome the inherent variations in the magnetoresistance ratio. For anything other than trivial cases, the state of the logic function must be determined by injecting a fixed current into the two strings of MTJs and measuring the difference in their terminal voltages with a differential amplifier as shown in Fig. 6.1. A fundamental limitation with this approach is that MTJs have no gain: the current passing through an MTJ is not sufficient to change the state of a subsequent MTJ.

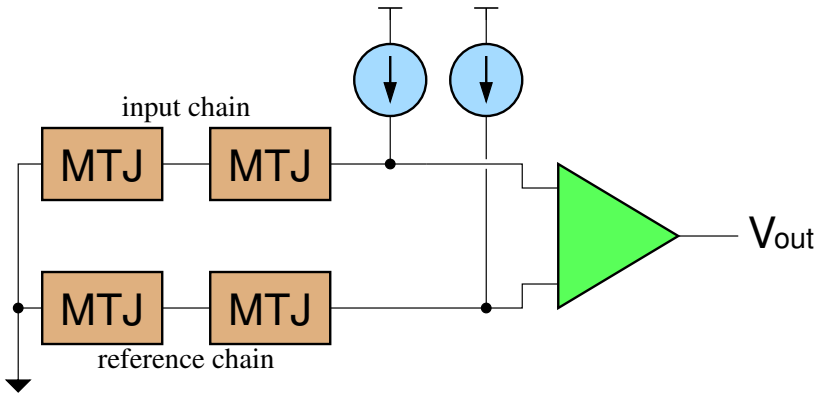

Figure 6.1: Programmable logic using MTJ chains

The final category of reconfigurable hardware using MTJs comprises those approaches that simply replace the configuration memory elements in the circuit with latches that use MTJs for non-volatile storage, such as those that are the subject of this research [162, 164, 174]. The MTJ latches require relatively high power consumption for the magnetic write operation, and the low magnetoresistance ratio implies a relatively slow read operation, so these latches are best used for configuring the interconnections and logical operations rather than acting as scratchpad memory for data processing [160]. Since the FPGA is itself a very regular structure, it may be advantageous to retain the XY crosspoint programming scheme used in

MRAMs, thereby reducing the overhead of incorporating the write current steering into every memory bit [161].

Chapter 7

Results

The foundation of this work required an understanding of several diverse topic areas. In order to understand the design constraints that necessarily resulted from using magnetic tunnel junctions an extensive review of their characteristics was undertaken. This included the micromagnetic behavior of thin ferromagnetic films and quantum tunneling. The bias dependence of the magnetoresistance ratio and the threat of dielectric breakdown were found to be the primary concerns when designing circuits around MTJs.

The generation of suitable magnetic fields for programming the MTJs was also a significant challenge. Simulations of the fields produced by currents in representative integrated circuit conductors verified that locating the MTJs against the conductor provided the most efficient generation of the necessary field strength, though the resulting current density was still quite high by conventional reliability standards. An in-depth review of electromigration phenomena suggested that this application could take advantage of the Blech effect, which allows short conducting paths to reliably support a much higher current density.

An evaluation of prior art circuit designs found that they were generally unsuitable for radiation-tolerant applications. Two new latch circuit designs were developed, a differential latch and a magnetic shadow latch. The differential latch utilized the MTJs to provide isolation between circuit nodes and thereby create redundant information storage. However, this circuit also had the disadvantage that there was no common node for the two MTJs that could serve as the connection to the write line, and there are concerns that single-event transients might damage the tunneling insulator.

The magnetic shadow latch takes advantage of the unique characteristics of a vertical

double-gate SOI technology, using the MTJs to control the bottom gate to preferentially unbalance the latch. The magnetic shadow latch concept was developed and optimized with typical circuit layouts, parasitic element extraction, and SPICE circuit simulation. The circuit design constraints were explored, resulting in a general understanding of the critical design parameters and trade-offs. The shadow latch was shown to be robust over wide variations in MTJ parameters, with several options for precharging the cell.

Finally, circuit techniques for implementing the one-wire programming scheme were explored. Detailed analyses explored the effect of parasitic resistance and evaluated the current density across the irregular shape of the write line. Various circuit configurations that reduced the average energy needed to program each latch were evaluated and optimized. As a result, it appears that magnetic memories can be programmed using much less energy than flash memories. However, the large currents required for programming will limit the use of magnetic memories to those applications that allow the simultaneous programming of a relatively small number of bits.

Of course, opportunities to expand this research remain. For example, the specific conductor patterns used to implement one-wire programming should be evaluated for electromigration reliability. Single-event effects testing should be conducted on typical MTJ structures, with and without a bias, to verify their resistance to damage from this type of radiation. More studies of the inherent reliability of Al_2O_3 and MgO barrier layers are needed.

In the area of circuit design, a more thorough characterization and modeling of the FlexFET bottom gate would give greater confidence in the circuit simulations. This would also allow the development of more efficient write-line drivers by using the bottom gate to temporarily lower the transistor threshold voltage. The possibility of extending the magnetic shadow latch concept to bulk CMOS deserves investigation.

Bibliography

[1] M. R. Ackermann and R. E. Mikawa. Factors contributing to CMOS static RAM upset. *IEEE Transactions on Nuclear Science*, NS-33(6):1524–1529, December 1986.

[2] William G. Moffatt, George W. Pearsall, and John Wulff. *The Structure and Properties of Materials, Volume I, Structure*. John Wiley and Sons, Inc., 1966.

[3] Gary A. Prinz. Spin-polarized transport. *Physics Today*, pages 58–63, April 1995.

[4] Gary A. Prinz. Magnetoelectronics. *Science*, 282:1660–1663, November 27, 1998.

[5] R. J. Soulen, Jr., J. M. Byers, M. S. Osofsky, B. Nadgorny, T. Ambrose, S. F. Cheng, P. R. Broussard, C. T. Tanaka, J. Nowak, J. S. Moodera, A. Barry, and J. M. D. Coey. Measuring the spin polarization of a metal with a superconducting point contact. *Science*, 282:85–88, October 2, 1998.

[6] F. J. Himpsel, J. E. Ortega, G. J. Mankey, and R. F. Willis. Magnetic nanostructures. *Advances in Physics*, 47(4):511–597, 1998.

[7] William Hart Hayt. *Engineering Electromagnetics*. McGraw-Hill, 3rd edition, 1974.

[8] Tera Analysis Ltd. Quickfield 5.4, October 24, 2006.

[9] T. Kishi, T. Ueda, H. Aikawa, T. Kajiyama, K. Shimura, Y. Fukuzumi, Y. Shimizu, Y. Asao, H. Yoda, and S. Tahara. Effect of novel yoke wire in low-power MRAM applications. *Journal of Magnetism and Magnetic Materials*, 272–276:1939–1940, 2004.

[10] Narain D. Arora, Li Song, Santosh Shah, Arani Sinha, and Victor Chang. Test chip for inductance characterization and modeling for sub-100nm X architecture and Manhattan chip design. In *International Conference on Microelectronic Test Structures*, volume 18, pages 251–255, April 2005.

[11] J. C. Fisher and I. Giaever. Tunneling through thin insulating layers. *Journal of Applied Physics*, 32(2):172–177, February 1961.

[12] John G. Simmons. Generalized formula for the electric tunnel effect between similar electrodes separated by a thin insulating film. *Journal of Applied Physics*, 34(6):1793–1803, June 1963.

[13] R. Meservey and P. M. Tedrow. Magnetic field splitting of the quasiparticle states in superconducting aluminum films. *Physical Review Letters*, 25(18):1270–1272, November 2, 1970.

[14] P. M. Tedrow and R. Meservey. Spin-dependent tunneling into ferromagnetic nickel. *Physical Review Letters*, 26(4):192–195, January 25, 1971.

[15] P. M. Tedrow and R. Meservey. Spin polarization of electrons tunneling from films of Fe, Co, Ni and Gd. *Physical Review B*, 7(1):318–326, 1973.

[16] M. Julliere. Tunneling between ferromagnetic films. *Physics Letters A*, 54(3):225–226, September 8, 1975.

[17] Terunobu Miyazaki and Nobuki Tezuka. Spin polarized tunneling in ferromagnet/insulator/ferromagnet junctions. *Journal of Magnetism and Magnetic Materials*, 151:403–410, 1995.

[18] J. S. Moodera, Lisa R. Kinder, Terrilyn M. Wong, and R. Meservey. Large magnetoresistance at room temperature in ferromagnetic thin film tunnel junctions. *Physical Review Letters*, 74(16):3273–3276, April 1995.

[19] T. Miyazaki and N. Tezuka. Giant magnetic tunneling effect in Fe/Al_2O_3/Fe junction. *Journal of Magnetism and Magnetic Materials*, 139:L231–L234, 1995.

[20] Evgeny Y. Tsymbal, Oleg N. Mryasov, and Patrick R. LeClair. Spin-dependent tunnelling in magnetic tunnel junctions. *Journal of Physics: Condensed Matter*, 15:R109–R142, February 5, 2003.

[21] R. Y. Gu, D. Y. Xing, and Jinming Dong. Spin-polarized tunneling between ferromagnetic films. *Journal of Applied Physics*, 80(12):7163–7165, December 15, 1996.

[22] J. M. Daughton. Magnetic tunneling applied to memory. *Journal of Applied Physics*, 81(8):3758–3763, April 15, 1997.

[23] Ricardo C. Sousa, Paulo P. Freitas, Virginia Chu, and João P. Conde. Vertical integration of a spin dependent tunnel junction with an amorphous Si diode for MRAM application. *IEEE Transactions on Magnetics*, 35(5):2832–2834, September 1999.

[24] S. Tehrani, J. M. Slaughter, E. Chen, M. Durlam, J. Shi, and M. DeHerrera. Progress and outlook for MRAM technology. *IEEE Transactions on Magnetics*, 35(5):2814–2819, September 1999.

[25] S. S. P. Parkin, K. P. Roche, M. G. Samant, P. M. Rice, , R. B. Beyers, R. E. Scheuerlein, E. J. O Sullivan, S. L. Brown, J. Bucchigano, D. W. Abraham, Yu Lu, M. Rooks, P. L. Trouilloud, R. A. Wanner, and W. J. Gallagher. Exchange-biased magnetic tunnel junctions and application to nonvolatile magnetic random access memory. *Journal of Applied Physics*, 85(8):5828–5833, April 15, 1999.

[26] S. Tehrani, B. Engel, J. M. Slaughter, E. Chen, M. DeHerrera, M. Durlam, P. Naji, R. Whig, J. Janesky, and J. Calder. Recent developments in magnetic tunnel junction MRAM. *IEEE Transactions on Magnetics*, 35(6):2752–2757, September 2000.

[27] Roy Scheuerlein, William Gallagher, Stuart Parkin, Alex Lee, Sam Ray, Ray Robertazzi, and William Reohr. A 10ns read and write non-volatile memory array

using a magnetic tunnel junction and FET switch in each cell. In *IEEE International Solid-State Circuits Conference*, pages 128–129, 2000.

[28] Stuart Parkin, Xin Jiang, Christian Kaiser, Alex Panchula, Kevin Roche, and Mahesh Samant. Magnetically engineered spintronic sensors and memory. *Proceedings of the IEEE*, 91(5):661–680, May 2003.

[29] Benedicte Warot, Amanda K. Petford-Long, and Thomas C. Anthony. Magnetic properties of patterned tunnel junctions. *Journal of Applied Physics*, 93(10):7287–7289, May 15, 2003.

[30] Peter Grünberg. Layered magnetic structures: History, highlights, applications. *Physics Today*, pages 31–37, May 2001.

[31] Yiming Huai, Frank Albert, Paul Nguyen, Mahendra Pakala, and Thierry Valet. Observation of spin-transfer switching in deep submicron-sized and low-resistance magnetic tunnel junctions. *Applied Physics Letters*, 84(16):3118–3120, April 19, 2004.

[32] Dexin Wang, Cathy Nordman, James M. Daughton, Zhenghong Qian, and Jonathon Fink. 70% TMR at room temperature for SDT sandwich junctions with CoFeB as free and reference layers. *IEEE Transactions on Magnetics*, 40(4):2269–2271, July 2004.

[33] Youfeng Zheng and Jian-Gang Zhu. Switching field variation in patterned submicron magnetic film elements. *Journal of Applied Physics*, 81(8):5471–5473, April 15, 1997.

[34] Ricardo C. Sousa and Paulo P. Freitas. Dynamic switching of tunnel junction MRAM cell with nanosecond field pulses. *IEEE Transactions on Magnetics*, 36(5):2770–2772, September 2000.

[35] Dan Wei, C. K. Ong, and Zheng Yang. Micromagnetic studies of read and write process in magnetoresistive random access memory. *Journal of Applied Physics*, 87(6):3068–3073, March 15, 2000.

[36] Ping Shang, Amanda K. Petford-Long, and Thomas C. Anthony. Effects of two in-plane fields on the magnetization reversal mechanism in magnetic tunnel junction elements. *Journal of Applied Physics*, 91(10):7703–7705, May 15, 2002.

[37] Y. K. Ha, J. E. Lee, H.-K. Kim, J. S. Bae, S. C. Oh, K. T. Nam, S. O. Park, N. I. Lee, H. K. Kang, U-In Chung, and M. T. Moon. MRAM with novel shaped cell using synthetic anti-ferromagnetic free layer. In *Symposium on VLSI Technology Digest of Technical Papers*, pages 24–25, 2004.

[38] Xiaochun Zhu and Jian-Gang Zhu. On the shape optimization of magnetic random access memory element design. *Journal of Applied Physics*, 93(10):8376–8378, May 15, 2003.

[39] C. C. Chang, Y. C. Chang, W. S. Chung, J. C. Wu, Z. H. Wei, M. F. Lai, and C. R. Chang. Influences of the aspect ratio and film thickness on switching properties of

elliptical permalloy elements. *IEEE Transactions on Magnetics*, 41(2):947–949, February 2005.

[40] Taewan Kim, Young Keun Kim, and Wanjun Park. Technological issues for high-density MRAM development. *Journal of Magnetism and Magnetic Materials*, 282:232–236, 2004.

[41] Taewan Kim, Sang Jin Park, Jinseo Noh, Wanjun Park, I-Hun Song, and Young Keun Kim. Current aspects and future perspectives of high-density MRAM. *Physica Status Solidi (A)*, 201(8):1617–1620, 2004.

[42] T. Mitsuzuka, K. Matsuda, A. Kamijo, and H. Tsuge. Interface structures and magnetoresistance in magnetic tunnel junctions. *Journal of Applied Physics*, 85(8):5807–5809, April 15, 1999.

[43] Th. Gerrits, H. A. M. van den Berg, J. Hohlfeld, L. Bär, and Th. Rasing. Ultrafast precessional magnetization reversal by picosecond magnetic field pulse shaping. *Nature*, 418:509–512, August 2002.

[44] Tzu-Ning Fang and Jian-Gang Zhu. 2D write addressability of tunneling junction MRAM elements. *IEEE Transactions on Magnetics*, 37(4):1963–1966, July 2001.

[45] Jyh-Shinn Yang, Ching-Rang Chang, W. C. Lin, and Denny D. Tang. Spin dynamics of unit cells in magnetic random access memory. *IEEE Transactions on Magnetics*, 41(2):879–882, February 2005.

[46] B. C. Choi, Q. F. Xiao, Y. K. Hong, S. H. Gee, J. Jabal, H. Han, K. J. Hass, and G. W. Donohoe. Numerical simulation study of magnetization precession dynamics in submicrometer elliptical $Ni_{80}Fe_{20}$ thin-film elements. *IEEE Transactions on Magnetics*, 42(10):3216–3218, October 2006.

[47] M. Motoyoshi, I. Yamamura, W. Ohtsuka, M. Shouji, H. Yamagishi, M. Nakamura, H. Yamada, K. Tai, T. Kikutani, T. Sagara, K. Moriyama, H. Mori, C. Fukumoto, M. Watanabe, H. Hachimo, H. Kano, K. Bessho, H. Narisawa, M. Hosomi, and N. Okazaki. A study for 0.18 μm high-density MRAM. In *Symposium on VLSI Technology Digest of Technical Papers*, pages 22–23, 2004.

[48] T. Kai, M. Yoshikawa, M. Nakayama, Y. Fukuzumi, T. Nagase, E. Kitagawa, T. Ueda, T. Kishi, S. Ikegawa, Y. Asao, K. Tsuchida, H. Yoda, N. Ishiwata, H. Hada, and S. Tahara. Improvement of robustness against write disturbance by novel cell design for high density MRAM. In *IEEE International Electron Devices Meeting Technical Digest*, pages 23.5.1–23.5.4, December 13-15, 2004.

[49] M. H. Park, Y. K. Hong, S. H. Gee, D. W. Erickson, and B. C. Choi. Magnetization configuration and switching behavior of submicron NiFe elements: Pac-man shape. *Applied Physics Letters*, 83(2):329–331, July 14, 2003.

[50] B. C. Choi, B. R. Pujada, Y. K. Hong, M. H. Park, H. Han, S. H. Gee, and G. W. Donohoe. Micromagnetic domain structures and magnetization switching mechanism in submicrometer thin-film elements. *IEEE Transactions on Magnetics*, 41(10):3109–3111, October 2005.

[51] B. R. Pujada, B. C. Choi, M. H. Park, Y. K. Hong, S. H. Gee, H. Han, and G. W. Donohoe. Micromagnetic configurations and switching mechanism in Pac-man-shaped submicron $Ni_{80}Fe_{20}$ magnets. *Journal of Applied Physics*, 97:073904, 2005.

[52] B. R. Pujada, J. Svendsen, K. O. Chipeniuk, B. C. Choi, M. H. Park, Y. K. Hong, S. H. Gee, and D. W. Erickson. Magnetic switching depending on as-patterned magnetization state in Pac-man shaped $Ni_{80}Fe_{20}$ submicron elements. *Journal of Applied Physics*, 96(8):4362–4365, October 15, 2004.

[53] M. H. Park, Y. K. Hong, S. H. Gee, D. W. Erickson, T. Tanaka, and B. C. Choi. Effect of shape anisotropy on switching behaviors of Pac-man NiFe submicron elements. *Journal of Applied Physics*, 95(11):7019–7021, June 1, 2004.

[54] B. C. Choi, J. Ho, Y. K. Hong, M. H. Park, H. Han, S. H. Gee, and G. W. Donohoe. Magnetization switching dynamics depending on as-patterned magnetization state in magnetic thin-film elements. *IEEE Transactions on Magnetics*, 41(10):2709–2711, October 2005.

[55] D. Meyners, H. Brückl, and G. Reiss. Influence of boundary roughness on the magnetization reversal in submicron sized magnetic tunnel junctions. *Journal of Applied Physics*, 93(5):2676–2680, March 1, 2003.

[56] Ching-Ray Chang and Jyh-Shinn Yang. The effect of structural defects on magnetic switching in thin ferromagnetic patterned films. *IEEE Transactions on Magnetics*, 43(2):923–926, February 2007.

[57] Ronald L. Pease. Total-dose issues for microelectronics in space systems. *IEEE Transactions on Nuclear Science*, 43(2):442–451, April 1996.

[58] P. S. Winokur, F. W. Sexton, G. L. Hash, and D. C. Turpin. Total-dose failure mechanisms of integrated circuits in laboratory and space environments. *IEEE Transactions on Nuclear Science*, NS-34(6):1448–1454, December 1987.

[59] J. V. Osborn, R. C. Lacoe, D. C. Mayer, and G. Yabiku. Total dose hardness of three commercial CMOS microelectronics foundries. *IEEE Transactions on Nuclear Science*, 45(3):1458–1463, June 1998.

[60] Ronald C. Lacoe, Jon V. Osborn, Rocky Koga, Stephanie Brown, and Donald C. Mayer. Application of hardness-by-design methodology to radiation-tolerant ASIC technologies. *IEEE Transactions on Nuclear Science*, 47(6):2334–2341, December 2000.

[61] G. Cellere, A. Paccagnella, A. Visconti, and M. Bonanomi. Ionizing radiation effects on floating gates. *Applied Physics Letters*, 85(3):485–487, July 19, 2004.

[62] Giorgio Cellere and Alessandro Paccagnella. A review of ionizing radiation effects in floating gate memories. *IEEE Transactions on Device and Materials Reliability*, 4(3):359–370, September 2004.

[63] Gabriel Molas, Damien Deleruyelle, Barbara De Salvo, Gérard Ghibaudo, Marc Gély, Luca Perniola, Dominique Lafond, and Simon Deleonibus. Degradation of floating-gate memory reliability by few electron phenomena. *IEEE Transactions on Electron Devices*, 53(10):2610–2619, October 2006.

[64] James C. Pickel and James T. Blanford, Jr. Cosmic-ray-induced errors in MOS devices. *IEEE Transactions on Nuclear Science*, 27(2):1006–1014, April 1980.

[65] Naval Research Laboratory. CREME96 [online]. July 30, 2004 [cited 20 Oct 2006]. Available from: `creme96.nrl.navy.mil`.

[66] E. Margaret Burbidge, G. R. Burbidge, William A. Fowler, and F. Hoyle. Synthesis of the elements in stars. *Reviews of Modern Physics*, 29(4):547–650, October 1957.

[67] Mieczyslaw Taube. *Evolution of Matter and Energy on a Cosmic and Planetary Scale*. Springer-Verlag, 1985.

[68] Brookhaven National Laboratory. TVDG LET calculator [online]. December 7, 2004 [cited January 2, 2007]. Available from: `tvdg10.phy.bnl.gov/LETCalc.html`.

[69] Richard L. Johnson, Jr., Sherra E. Diehl-Nagle, and John R. Hauser. Simulation approach for modeling single event upsets on advanced CMOS SRAMs. *IEEE Transactions on Nuclear Science*, NS-32(6):4122–4127, December 1985.

[70] John R. Hauser, Sherra E. Diehl-Nagle, A. R. Knudson, A. B. Campbell, W. J. Stapor, and P. Shapiro. Ion track shunt effects in multi-junction structures. *IEEE Transactions on Nuclear Science*, NS-32(6):4115–4121, December 1985.

[71] J. P. Kreskovsky and H. L. Grubin. Simulation of charge collection in a multilayer device. *IEEE Transactions on Nuclear Science*, NS-32(6):4140–4144, December 1985.

[72] Alexander O. Brown, Bharat Bhuva, S. E. Kerns, and W. J. Stapor. Practical approach to determining charge collected in multi-junction structures due to the ion shunt effect. *IEEE Transactions on Nuclear Science*, 40(6):1918–1925, December 1993.

[73] H. Dussault, J. W. Howard, Jr., R. C. Block, M. R. Pinto, W. J. Stapor, and A. R. Knudson. High energy heavy-ion-induced single event transients in epitaxial structures. *IEEE Transactions on Nuclear Science*, 41(6):2018–2025, December 1994.

[74] E. L. Petersen, J. C. Pickel, J. H. Adams, Jr., and E. C. Smith. Rate prediction for single event effects. *IEEE Transactions on Nuclear Science*, 39(6):1577–1599, December 1992.

[75] P. E. Dodd, F. W. Sexton, and P. S. Winokur. Three-dimensional simulation of charge collection and multiple-bit upset in Si devices. *IEEE Transactions on Nuclear Science*, 41(6):2005–2017, December 1994.

[76] O. Musseau. Single-event effects in SOI technologies and devices. *IEEE Transactions on Nuclear Science*, 43(2):603–703, April 1996.

[77] J. R. Schwank, V. Ferlet-Cavrois, M. R. Shaneyfelt, P. Paillet, and P. E. Dodd. Radiation effects in SOI technologies. *IEEE Transactions on Nuclear Science*, 50(3):522–538, June 2003.

[78] V. Ferlet-Cavrois, G. Gasiot, C. Marcandella, C. DHose, O. Flament, O. Faynot, J. du Port de Pontcharra, and C. Raynaud. Insights on the transient response of fully and partially depleted SOI technologies under heavy-ion and dose-rate irradiations. *IEEE Transactions on Nuclear Science*, 49(6):2948–2956, December 2002.

[79] Eric P. Ver Ploeg, Cuong T. Nguyen, S. Simon Wong, and James D. Hummer. Parasitic bipolar gain in fully depleted n-channel SOI MOSFET's. *IEEE Transactions on Electron Devices*, 41(6):970–977, June 1994.

[80] G. Bruguier and J-M. Palau. Single particle-induced latchup. *IEEE Transactions on Nuclear Science*, 43(2):522–532, April 1996.

[81] A. H. Johnston. The influence of VLSI technology evolution on radiation-induced latchup in space systems. *IEEE Transactions on Nuclear Science*, 43(2):505–521, April 1996.

[82] Roberto Menozzi, Luca Selmi, Enrico Sangiorgi, Guiseppe Crisenza, Tiziana Cavioni, and Bruno Ricco. Layout dependence of CMOS latchup. *IEEE Transactions on Electron Devices*, 35(11):1892–1901, November 1988.

[83] G. Cellere, A. Paccagnella, A. Visconti, and M. Bonanomi. Soft Errors induced by single heavy ions in Floating Gate memory arrays. In *20th IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems*, 2005.

[84] G. Cellere, A. Paccagnella, A. Visconti, M. Bonanomi, and A. Candelori. Transient conductive path induced by a single ion in 10 nm SiO2 layers. *IEEE Transactions on Nuclear Science*, 51(6):3304–3311, December 2004.

[85] L. W. Massengill, B. K. Choi, D. M. Fleetwood, R. D. Schrimpf, K. F. Galloway, M. R. Shaneyfelt, T. L. Meisenheimer, P. E. Dodd, J. R. Schwank, Y. M. Lee, R. S. Johnson, and G. Lucovsky. Heavy-ion-induced breakdown in ultra-thin gate oxides and high-k dielectrics. *IEEE Transactions on Nuclear Science*, 48(6):1904–1912, December 2001.

[86] B. K Choi, D. M. Fleetwood, R. D. Schrimpf, L. W. Massengill, K. F. Galloway, M. R. Shaneyfelt, T. L. Meisenheimer, P. E. Dodd, J. R. Schwank, Y. M. Lee, R. S. John, and G. Lucovsky. Long-term reliability degradation of ultrathin dielectric films due to heavy-ion irradiation. *IEEE Transactions on Nuclear Science*, 49(6):3045–3050, December 2002.

[87] Gary K. Lum, Nicholas Boruta, J. M. Baker, L. Robinette, M. R. Shaneyfelt, J. R. Schwank, P. E. Dodd, and J. A. Felix. New experimental findings for single-event gate rupture in MOS capacitors and linear devices. *IEEE Transactions on Nuclear Science*, 51(6):3263–3269, December 2004.

[88] H. T. Weaver, C. L. Axness, J. D. McBrayer, J. S. Browning, J. S. Fu, A. Ochoa, Jr., and R. Koga. An SEU tolerant memory cell derived from fundamental studies of SEU mechanisms in SRAM. *IEEE Transactions on Nuclear Science*, NS-34(6):1281–1286, December 1987.

[89] F. W. Sexton, R. K. Treece, K. J. Hass, K. L. Hughes, G. L. Hash, C. L. Axness, S. P. Buchner, and K. Kang. SEU characterization and design dependence of the SA3300 microprocessor. *IEEE Transactions on Nuclear Science*, 37(6):1861–1868, December 1990.

[90] F. W. Sexton, W. T. Corbett, R. K. Treece, K. J. Hass, K. L. Hughes, C. L. Axness, G. L. Hash, M. R. Shaneyfelt, and T. F. Wunsch. SEU simulation and testing of resistor-hardened D-latches in the SA3300 microprocessor. *IEEE Transactions on Nuclear Science*, 38(6):1521–1528, December 1991.

[91] Matthew J. Gadlage, Ronald D. Schrimpf, Joseph M. Benedetto, Paul H. Eaton, David G. Mavis, Mike Sibley, Keith Avery, and Thomas L. Turflinger. Single event transient pulsewidths in digital microcircuits. *IEEE Transactions on Nuclear Science*, 51(6):3285–3290, December 2004.

[92] P. E. Dodd, F. W. Sexton, G. L. Hash, M. R. Shaneyfelt, B. L. Draper, A. J. Farino, and R. S. Flores. Impact of technology trends on SEU in CMOS SRAMs. *IEEE Transactions on Nuclear Science*, 43(6):2797–2804, December 1996.

[93] Leonard R. Rockett. Designing CMOS data cells for space systems. *Microelectronics Journal*, 35:953–967, 2004.

[94] Leonard R. Rockett, Jr. An SEU-hardened CMOS data latch design. *IEEE Transactions on Nuclear Science*, 35(6):1682–1687, December 1988.

[95] T. Calin, M. Nicolaidis, and R. Velazco. Upset hardened memory design for submicron CMOS technology. *IEEE Transactions on Nuclear Science*, 43(6):2874–2878, December 1996.

[96] P. Hazucha, T. Karnik, S. Walstra, B. A. Bloechel, J. W. Tschanz, J. Maiz, K. Soumyanath, G. E. Dermer, S. Narendra, V. De, and S. Borkar. Measurements and analysis of SER-tolerant latch in a 90-nm dual-V_T CMOS process. *IEEE Journal of Solid-State Circuits*, 39(9):1536–1543, September 2004.

[97] Michael Nicolaidis. Design for soft error mitigation. *IEEE Transactions on Device and Materials Reliability*, 5(3):405–418, September 2005.

[98] J. Gambles, L. Miles, J. Hass, W. Smith, S. Whitaker, and B. Smith. An ultra-low-power, radiation-tolerant Reed Solomon encoder for space applications. In *IEEE Custom Integrated Circuits Conference*, pages 631–634, September 21–24, 2003.

[99] Jody W. Gambles, Kenneth J. Hass, and Sterling R. Whitaker. Radiation hardness of ultra low power CMOS VLSI. In *11th NASA Symposium on VLSI Design*, May 2003.

[100] Michael J. Barry. Radiation resistant SRAM memory cell. *U.S. Patent 5,157,625*, October 20, 1992.

[101] Jerry G. Dooley. SEU-immune latch for gate array, standard cell, and other ASIC applications. *U.S. Patent 5,311,070*, May 10, 1994.

[102] A. Makihara, T. Yamaguchi, Y. Tsuchiya, T. Arimitsu, H. Asai, Y. Iide, H. Shindou, S. Kuboyama, and S. Matsuda. SEE in a 0.15 μm fully depleted CMOS/SOI commercial process. *IEEE Transactions on Nuclear Science*, 51(6):3621–3625, December 2004.

[103] Prashant Kumar Saxena and Navakanta Bhat. SEU reliability improvement due to source-side charge collection in the deep-submicron SRAM cell. *IEEE Transactions on Device and Materials Reliability*, 3(1):14–17, March 2003.

[104] Jeffrey D. Black, Andrew L. Sternberg, Michael L. Alles, Arthur F. Witulski, Bharat L. Bhuva, Lloyd W. Massengill, Joseph M. Benedetto, Mark P. Baze, Jerry L. Wert, and Matthew G. Hubert. HBD layout isolation techniques for multiple node charge collection mitigation. *IEEE Transactions on Nuclear Science*, 52(6):2536–2541, December 2005.

[105] S. T. Chui. Bias dependence in spin-polarized tunneling. *Physical Review B*, 55(9):5600–5603, March 1, 1997.

[106] Yu Lu, X. W. Li, Gang Xiao, R. A. Altman, W. J. Gallagher, A. Marley, K. Roche, and S. Parkin. Bias voltage and temperature dependence of magnetotunneling effect. *Journal of Applied Physics*, 83(11):6515–6517, June 1, 1998.

[107] R. C. Sousa, J. J. Sun, V. Soares, P. P. Freitas, A. Kling, M. F. da Silva, and J. C. Soares. Temperature dependence and annealing effects on spin dependent tunnel junctions. *Journal of Applied Physics*, 85(8):5258–5260, April 15, 1999.

[108] J. Zhang and R. M. White. Voltage dependence of magnetoresistance in spin dependent tunneling junctions. *Journal of Applied Physics*, 83(11):6512–6514, June 1, 1998.

[109] Chang He Shang, Janusz Nowak, Ronnie Jansen, and Jagadeesh S. Moodera. Temperature dependence of magnetoresistance and surface magnetization in ferromagnetic tunnel junctions. *Physical Review B*, 58(6):R2917–R2920, August 1, 1998.

[110] T. Hagler, R. Kinder, and G. Bayreuther. Temperature dependence of tunnel magnetoresistance. *Journal of Applied Physics*, 89(11):7570–7572, June 1, 2001.

[111] Jagadeesh S. Moodera, Janusz Nowak, and Rene J. M. van de Veerdonk. Interface magnetism and spin wave scattering in ferromagnet-insulator-ferromagnet tunnel junctions. *Physical Review Letters*, 80(13):2941–2944, March 30, 1998.

[112] X. Kou, J. Schmalhorst, A. Thomas, and G. Reiss. Temperature dependence of the resistance of magnetic tunnel junctions with MgO barrier. *Applied Physics Letters*, 88(21):212115, May 2006.

[113] A. H. Davis, J. M. MacLaren, and P. LeClair. Inherent temperature effects in magnetic tunnel junctions. *Journal of Applied Physics*, 89(11):7567–7569, June 1, 2001.

[114] X. H. Xiang, T. Zhu, J. Du, G. Landry, and John Q. Xiao. Effect of density of states on bias dependence in magnetic tunnel junctions. *Physical Review B*, 66:174407, 2002.

[115] Jerzy Krupka, Richard G. Geyer, Matthias Kuhn, and Johann Heyen Hinken. Dielectric properties of single crystals of Al_2O_3, $LaAlO_3$, $NdGaO_3$, $SrTiO_3$ and MgO at cryogenic temperatures. *IEEE Transactions on Microwave Theory and Techniques*, 42(10):1886–1889, October 1994.

[116] J. E. Aitken, P. H. Ladbrooke, and M. H. N. Potok. Microwave measurement of the temperature coefficient of permittivity for sapphire and alumina. *IEEE Transactions on Microwave Theory and Techniques*, pages 526–529, June 1975.

[117] G. J. Sloggett, N. G. Barton, and S. J. Spencer. Fringing fields in disc capacitors. *Journal of Physics A: Mathematical and General*, 19(14):2725–2736, 1986.

[118] T. Sakurai and K. Tamaru. Simple formulas for two- and three-dimensional capacitances. *IEEE Transactions on Electron Devices*, ED-30(2):183–185, 1983.

[119] H. Kaiju, S. Fujita, T. Morozumi, and K. Shiiki. Magnetocapacitance effect of spin tunneling junctions. *Journal of Applied Physics*, 91(10):7430–7432, May 15, 2002.

[120] J. A. Felix, J. R. Schwank, D. M. Fleetwood, M. R. Shaneyfelt, and E. P. Gusev. Effects of radiation and charge trapping on the reliability of high-κ gate dielectrics. *Microelectronics Reliability*, 44:563–575, 2004.

[121] Y. Conraux, J. P. Nozières, V. Da Costa, M. Toulemonde, and K. Ounadjela. Effects of swift heavy ion bombardment on magnetic tunnel junction functional properties. *Journal of Applied Physics*, 93(10):7301–7303, May 15, 2003.

[122] W. Oepts, H. J. Verhagen, W. J. M. de Jonge, and R. Coehoorn. Dielectric breakdown of ferromagnetic tunnel junctions. *Applied Physics Letters*, 73(16):2363–2365, October 19, 1998.

[123] W. Oepts, H. J. Verhagen, R. Coehoorn, and W. J. M. de Jonge. Analysis of breakdown in ferromagnetic tunnel junctions. *Journal of Applied Physics*, 86(7):3863–3872, October 1, 1999.

[124] K. Shimazawa, N. Kasahara, J. J. Sun, S. Araki, H. Morita, and M. Matsuzaki. Electrical breakdown of the magnetic tunneling junction with an AlO_x barrier formed by radical oxidation. *Journal of Applied Physics*, 87(9):5194–5196, May 2000.

[125] J. Das, R. Degraeve, H. Boeve, P. Duchamps, L. Lagas, G. Groeseneken, G. Borghs, and J. De Boeck. Degradation and time dependent breakdown of stressed ferromagnetic tunnel junctions. *Journal of Applied Physics*, 89(11):7350–7352, June 1, 2001.

[126] D. Rao, K. Sin, M. Gibbons, S. Funada, M. Mao, C. Chien, and H.-C. Tong. Voltage-induced barrier-layer damage in spin-dependent tunneling junctions. *Journal of Applied Physics*, 89(11):7362–7364, June 1, 2001.

[127] Chikako Yoshida, Hideyuki Noshiro, Yuichi Yamazaki, Takashi Iizuka, Yoshihiro Stoh, Masaki Aoki, Shinnjiro Umehara, Masashige Satoh, and Kazuo Kobayashi. Reliability of MgO tunneling barrier for MRAM device. In *International Reliability Physics Symposium*, pages 697–698, 2006.

[128] D. G. Pierce and P. G. Brusius. Electromigration: A review. *Microelectronics Reliability*, 37(7):1053–1072, 1997.

[129] Christine S. Hau-Riege. An introduction to Cu electromigration. *Microelectronics Reliability*, 44:195–205, 2004.

[130] Changsup Ryu, Kee-Won Kwon, Alvin L. S. Loke, Haebum Lee, Takeshi Nogami, Valery M. Dubin, Rahim A. Kavari, Gary W. Ray, and S. Simon Wong. Microstructure and reliability of copper interconnects. *IEEE Transactions on Electron Devices*, 46(6):1113–1120, June 2005.

[131] Jiang Tao, Nathan W. Cheung, and Chenming Hu. Electromigration characteristics of copper interconnects. *IEEE Transactions on Electron Devices*, 14(5):249–251, May 1993.

[132] Paul Besser, Amit Marathe, Larry Zhao, Matthew Herrick, Cristiano Capasso, and Hisao Kawasaki. Optimizing the electromigration performance of copper interconnects. In *IEEE International Electron Devices Meeting Technical Digest*, pages 6.1.1–6.1.4, 2000.

[133] I. A. Blech. Electromigration in thin aluminum films on titanium oxide. *Journal of Applied Physics*, 47(4):1203–1208, April 1976.

[134] E. M. Atakov, T. S. Sriram, D. Dunnell, and S. Pizzanello. Effect of VLSI interconnect layout on electromigration performance. In *36th Annual International Reliability Physics Symposium*, pages 348–355, 1998.

[135] R. A. Wachnik, R. G. Filippi, T. M. Shaw, and P. C. Lin. Practical benefits of the electromigration short-length effect, including a new design rule methodology and an electromigration resistant power grid with enhanced wireability. In *Symposium on VLSI Technology Digest of Technical Papers*, pages 220–221, 2000.

[136] Ennis T. Ogawa, Alexander J. Bierwag, Ki-Don Lee, Hideki Matsuhashi, Patrick R. Justison, Anup N. Ramamurthi, Paul S. Hoa, Volker A. Blaschke, David Griffiths, Anne Nelsen, Mark Breen, and Robert H. Havemann. Direct observation of a critical length effect in dual-damascene Cu/oxide interconnects. *Applied Physics Letters*, 78(18):2652–2654, April 30, 2005.

[137] Joris Proost, Luc Delaey, Jan D'Haen, and Karen Maex. Plasticity of electromigration-induced hillocking and its effect on the critical length. *Journal of Applied Physics*, 91(11):9108–9115, June 1, 2002.

[138] M. H. Lin, K. P. Chang, K. C. Su, and Tahui Wang. Effects of width scaling, length scaling, and layout variation on electromigration in dual damascene copper interconnects. In *International Reliability Physics Symposium*, pages 671–672, 2006.

[139] Shinji Yokogawa and Hideaki Tsuchiya. Effects of Al doping on the electromigration performance of damascene Cu interconnects. *Journal of Applied Physics*, 101:013513, 2007.

[140] Richard Frankovic, Gary H. Bernstein, and J. Joseph Clement. Pulsed-current duty cycle dependence of electromigration-induced stress generation in aluminum conductors. *IEEE Electron Device Letters*, 17(5):244–246, May 1996.

[141] Xiang Gui, James W. Haslett, Steven K. Dew, and Michael J. Brett. Simulation of temperature cycling effects on electromigration behavior under pulsed current stress. *IEEE Transactions on Electron Devices*, 45(2):380–386, February 1998.

[142] James E. Murguia and Joseph B. Bernstein. Short-time failure of metal interconnect caused by current pulses. *IEEE Electron Device Letters*, 14(10):481–483, October 1993.

[143] Xiang Gui, Steven K. Dew, and Michael J. Brett. Thermal simulation of thin-film interconnect failure caused by high current pulses. *IEEE Transactions on Electron Devices*, 42(7):1386–1388, July 1995.

[144] Kaustav Banerjee, Ajith Amerasekera, Nathan Cheung, and Chenming Hu. High-current failure model for VLSI interconnects under short-pulse conditions. *IEEE Electron Device Letters*, 18(9):405–407, September 1997.

[145] John Scarpulla, David C. Eng, Stephanie Brown, and Kenneth P. MacWilliams. Reliability of metal interconnect after a high-current pulse. *IEEE Electron Device Letters*, 17(7):322–324, July 1996.

[146] N. Tezuka, M. Oogane, and T. Miyazaki. Applied voltage and temperature dependence of tunneling magnetoresistance. *Journal of Magnetism and Magnetic Materials*, pages 149–151, 1999.

[147] Saied Tehrani, Jon M. Slaughter, Mark DeHerrera, Brad N. Engel, Nicholas D. Rizzo, John Salter, Mark Durlam, Renu W. Dave, Jason Janesky, Brian Butcher, Ken Smith,

and Greg Grynkewich. Magnetoresistive random access memory using magnetic tunnel junctions. *Proceedings of the IEEE*, 91(5):703–712, May 2003.

[148] Jagadeesh S. Moodera and George Mathon. Spin polarized tunneling in ferromagnetic junctions. *Journal of Magnetism and Magnetic Materials*, pages 248–273, 1999.

[149] Edward S. Yang. *Fundamentals of Semiconductor Devices*. McGraw-Hill, 1978.

[150] Bodhisattva Das and William C. Black, Jr. A generalized HSPICE macro-model for pinned spin-dependent-tunneling devices. *IEEE Transactions on Magnetics*, 35(5):2889–2891, September 1999.

[151] Chung-Hyun Cho, Ju Hyon Ko, and Daejeong Kim. A CMOS macro-model for MTJ resistor of MRAM cell. *Physica Status Solidi (A)*, 201(8):1653–1657, 2004.

[152] A. Bette, J. DeBrosse, D. Gogl, H. Hoenigschmid, R. Robertazzi, C. Arndt, D. Braun, D. Casarotto, R. Havreluk, S. Lammers, W. Obermaier, W. Reohr, H. Viehmann, W. J. Gallagher, and G. Müller. A high-speed 128kbit MRAM core for future universal memory applications. In *Symposium on VLSI Circuits Digest of Technical Papers*, 2003.

[153] Edward K. S. Au, Wing-Hung Ki, Wai Ho Mow, Silas T. Hung, and Catherine Y. Wong. A novel current-mode sensing scheme for magnetic tunnel junction RAM. *IEEE Transactions on Magnetics*, 40(2):483–488, March 2004.

[154] Kouichi Yamada, Naofumi Sakai, Yoshiyuki Ishizuka, and Kazunobu Mameno. A novel sensing scheme for a MRAM with a 5% MR ratio. In *Symposium on VLSI Circuits Digest of Technical Papers*, pages 123–124, 2001.

[155] T. Inaba, K. Tsuchida, T. Sugibayashi, S. Tahara, and H. Yoda. Resistance ratio read (R^3) architecture for a burst operated 1.5V MRAM macro. In *IEEE Custom Integrated Circuits Conference*, pages 399–402, 2003.

[156] Masaki Aoki, Hiroshi Iwasa, and Yoshihiro Sato. A novel voltage sensing 1T/2MTJ cell with resistance ratio for highly stable and scalable MRAM. In *Symposium on VLSI Technology Digest of Technical Papers*, pages 170–171, 2005.

[157] James M. Daughton, Brenda A. Everitt, and Arthur V. Pohm. Spin dependent tunneling memory. *U.S. Patent 6,275,411*, August 14, 2001.

[158] D. Wang, M. Tondra, A. V. Pohm, C. Nordman, J. Anderson, and J. M. Daughton. Spin dependent tunneling devices fabricated for magnetic random access memory applications using latching mode. *Journal of Applied Physics*, 87(9):6385–6387, May 1, 2000.

[159] William C. Black, Jr. and Bodhisattva Das. Programmable logic using giant-magnetoresistance and spin-dependent tunneling devices. *Journal of Applied Physics*, 87(9):6674–6679, May 1, 2000.

[160] W. Zhao, E. Belhaire, V. Javerliac, C. Chappert, and B. Dieny. Evaluation of a non-volatile FPGA based on MRAM technology. In *IEEE International Conference on Integrated Circuit Design and Technology*, pages 1–4, May 24-26, 2006.

[161] N. Bruchon, L. Torres, G. Sassatelli, and G. Cambon. New non-volatile FPGA concept using magnetic tunneling junction. In *Emerging VLSI Technologies and Architectures*, pages 269–276, March 2–3, 2006.

[162] K. J. Hass, G. W. Donohoe, and Y. K. Hong. SEU-resistant magnetic flip flops. In *12th NASA VLSI Symposium*, October 2005.

[163] K. J. Hass, G. W. Donohoe, Y.-K. Hong, and B. C. Choi. Magnetic latches for space applications. In *IEEE International Magnetics Conference*, May 8-12, 2006.

[164] K. J. Hass, G. W. Donohoe, Y.-K. Hong, and B. C. Choi. Magnetic flip flops for space applications. *IEEE Transactions on Magnetics*, 42(10):2751–2753, October 2006.

[165] Kenneth J. Hass, Gregory Donohoe, Yang-Ki Hong, Byoung-Chul Choi, Kelly DeGregorio, and Richard Hayhurst. Integrated magnetic memory for embedded computing systems. In *IEEE Aerospace Conference*, pages 1–10, March 5–9, 2007.

[166] John S. Bialas, Jr. and Joseph A. Hoffman. Single event upset immune register with fast write access. *U.S. Patent 5,525,923*, June 11, 1996.

[167] S. Parke, M. Goldston, D. Hackler, K. DeGregorio, R. Hayhurst, A. Horvath, and S. Parsa. Complementary sidewall-spacer-diffused ultrashallow SD extension process for damascene independently-double-gated SOI CMOS. In *IEEE International SOI Conference*, pages 104–105, October 4–7, 2004.

[168] S. Parke, K. DeGregorio, M. Goldston, R. Hayhurst, and D. Hackler. Flexfet: A low-cost, rad-hard, independent-double-gate SOI CMOS technology with flexible, dynamic reconfigurability. In *IEEE Aerospace Conference*, pages 1–8, March 5–12, 2005.

[169] S. Parke, K. DeGregorio, D. Hackler, and D. Wilson. Ultra-low-power dynamic-threshold digital circuits in the flexFET independently-double-gated SOI CMOS technology. In *IEEE Workshop on Microelectronics and Electron Devices*, pages 33–34, April 14, 2006.

[170] S. Parke, M. Goldston, D. Hackler, K. DeGregorio, R. Hayhurst, A. Horvath, and S. Parsa. FlexFET: a low-cost planar, self-aligned independent-double-gate MOSFET with flexible, dynamic threshold control. In *IEEE Workshop on Microelectronics and Electron Devices*, pages 35–37, April 15, 2005.

[171] Kaushik Roy, Hamid Mahmoodi, Saibal Mukhopadhyay, Hari Ananthan, Aditya Bansai, and Tamer Cakici. Double-gate SOI devices for low-power and high-performance applications. In *IEEE/ACM International Conference on Computer Aided Design*, pages 217–224, November 6–10, 2005.

[172] D. Hackler, K. DeGregorio, and S. Parke. Ultra-low-power, high-performance, dynamic-threshold digital circuits in the FlexFET independently-double-gated SOI CMOS technology. In *IEEE International SOI Conference*, pages 81–82, October 3–6, 2005.

[173] S. Parke. Comparison of existing & proposed SOI MOSFET device structures for minimizing total dose radiation damage. In *IEEE Aerospace Conference*, volume 4, pages 2427–2430, March 6–13, 2004.

[174] K. J. Hass, G. W. Donohoe, Y.-K. Hong, B.-C. Choi, K. DeGregorio, and R. Hayhurst. Magnetic shadow RAM. In *Non-Volatile Memory Technology Symposium*, November 5–8, 2006.

[175] Jonathan E. Knudsen and Lawrence T. Clark. An area and power efficient radiation hardened by design flip-flop. *IEEE Transactions on Nuclear Science*, 53(6):3392–3399, December 2006.

[176] Brad N. Engel, Nicholas D. Rizzo, Jason Janesky, Jon M. Slaughter, Renu Dave, Mark DeHerrera, Mark Durlam, and Saied Tehrani. The science and technology of magnetoresistive tunneling memory. *IEEE Transactions on Nanotechnology*, 1(1):32–38, March 2002.

[177] B. C. Choi and Y. K. Hong. Private communication, October 11, 2005.

[178] Mark Durlam, Peter J. Naji, Asim Omair, Mark DeHerrera, John Calder, Jon M. Slaughter, Brad N. Engel, Nicholas D. Rizzo, Greg Grynkewich, Brian Butcher, Clarence Tracy, Ken Smith, Kelly W. Kyler, J. Jack Ren, Jaynal A. Molla, William A. Feil, Rick G. Williams, and Saiad Tehrani. A 1-Mbit MRAM based on 1T1MTJ bit cell integrated with copper interconnects. *IEEE Journal of Solid-State Circuits*, 38(5):769–773, May 2003.

[179] Y. Asao, T. Kajiyama, Y. Fukuzumi, M. Amano, H. Aikawa, T. Ueda, T. Kishi, S. Ikegawa, K. Tsuchida, Y. Iwata, A. Nitayama, K. Shimura, Y. Kato, S. Miura, N. Ishiwata, H. Hada, S. Tahara, and H. Yoda. Design and process integration for high-density, high-speed, and low-power $6F^2$ cross point MRAM cell. In *IEEE International Electron Devices Meeting Technical Digest*, pages 23.2.1–23.2.4, December 13-15, 2004.

[180] Hyung Gyu Lee and Naehyuck Chang. Energy-aware memory allocation in heterogeneous non-volatile memory systems. In *International Symposium on Low Power Electronics and Design*, pages 420–423, August 25–27, 2003.

[181] Gaurav Mathur, P. Desnoyers, Deepak Ganesan, and Prashant Shenoy. Ultra-low power data storage for sensor networks. In *The Fifth International Conference on Information Processing in Sensor Networks*, pages 374–381, April 19–21, 2006.

[182] Chenming Hu. Interconnect devices for field programmable gate array. In *International Electron Devices Meeting*, pages 24.1.1–24.1.4, 1992.

[183] A. Ney, C. Pampuch, R. Koch, and K. H. Ploog. Programmable computing with a single magnetoresistive element. *Nature*, 425:485–487, October 2, 2003.

[184] A. Ney and J. S. Harris, Jr. Reconfigurable magnetologic computing using the spin flop switching of a magnetic random access memory cell. *Applied Physics Letters*, 86:013502, 2005.

[185] G. Reiss, H. Brückl, A. Hütten, H. Koop, D. Meyners, A. Thomas, S. Kämmerer, J. Schmalhorst, and M. Brzeska. New materials and applications for magnetic tunneling junctions. *Physica Status Solidi (A)*, 201(8):1628–1634, 2004.

[186] D. Meyners, K. Rott, H. Brückl, G. Reiss, and J. Wecker. Submicron-sized magnetic tunnel junctions in field programmable logic gate arrays. *Journal of Applied Physics*, 99:023907, 2006.

Appendices

Appendix A

Magnetic Field Software Model

The program listing below calculates the magnetic field produced in some "area of interest" near a rectangular conductor (wire). All materials are assumed to have a relative permeability of 1.0 and to be isotropic. The program is written in *python*, an interpreted, object-oriented language. A unique aspect of python is that leading white space in a line is significant, and indentation is used to create sections of code corresponding to conditional and looping constructs.

The first ten lines of the program are standard boiler-plate for python and are not specific to the problem at hand. In line 13 the current flowing through the wire (in the Z dimension) is specified, in amperes. In this case the current is 1 mA.

The parameters used to define the physical dimensions of the wire and the area of interest are illustrated in Fig. A.1. The rectangle labelled `WIRE` represents the rectangular cross section of the conductor, with current flowing into of the plane of the paper. Although the origin of the coordinates in this model can be at any arbitrary location with respect to the wire and the area of interest, the data shown in Section 1.3 assumes that the origin is on the upper edge of the wire, halfway between its left and right edges.

Lines 15 through 18 define the size and location of the rectangular cross section of the wire. This example shows a wire that is 1 μm wide (in the X dimension) and 2 μm thick (in the Y dimension). Lines 20 through 23 define the area of interest. The magnetic field will be computed only for points within this rectangle, which is 1 μm wide but only 1 nm high in this example.

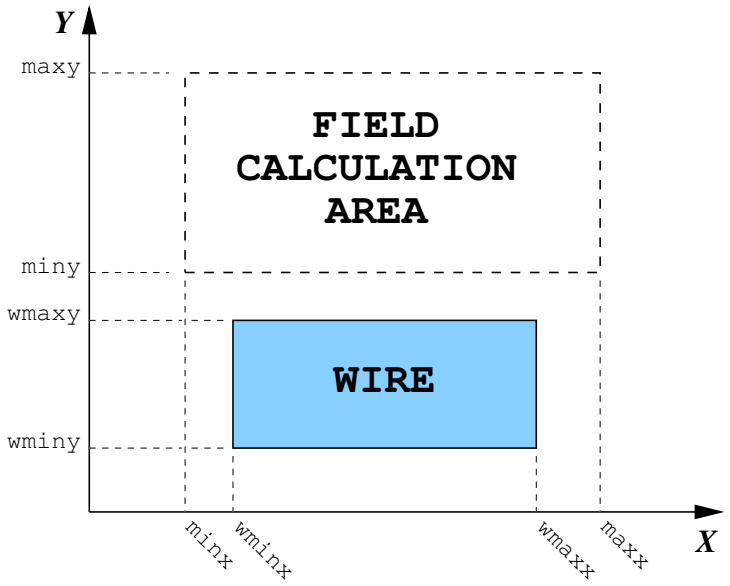

Figure A.1: Layout parameters for software model

Line 25 defines the grid size for the model. The cross section of the wire is subdivided using this grid, and the conductor in each grid cell is modeled as an infinitesimally thin, infinitely long wire located at the center of the cell. The current flowing through each of these ideal wires is calculated in line 27, but dividing the total current by the total number of grid squares in the wire. In line 28 the magnetic field strength produced at a distance of 1 meter from one of the ideal wires is calculated, in units of A/m, using

$$H = \frac{\mu i}{2\pi r} \tag{A.1}$$

This value is converted to oersteds in line 29, where

$$1 \text{ Oe} = 1 \text{ A/m} \times \frac{1000}{4\pi} \approx 79.6 \text{ A/m} \tag{A.2}$$

The input parameters and the calculated current density are printed to standard output in lines 31 through 33. Standard output is typically redirected to a file, so this provides some documentation of the input conditions that generated the data to follow.

The program keeps track of the minimum and maximum values of the X and Y components of the field over all points in the field of interest. The variables that will hold these values are initialized to dummy values in lines 35 through 38.

In line 39 the value of one-half the grid size is calculated. This is the distance from the edges of each cell to the center of the cell, whether in the conductor or in the area of interest.

The real work begins in line 41, where the first X coordinate in the area of interest is selected. Line 42 then causes lines 43 through 67 to iterate over all of the X coordinates for cells in the area of interest. Similarly, line 43 selects the initial Y coordinate, and line 44 causes the program to iterate lines 45 through 67 to iterate over all of the Y coordinates for cells in the area of interest.

The calculation of the field at a given point in the area of interest begins in lines 45 and 46. The field contribution from each cell in the conductor will be summed, so the variables that will hold these sums (`xv` and `yv`) must be cleared. The calculations of the field components have been optimized for speed in the software and are not obvious when reading the code, but can be explained with the help of Fig. A.2. Note that `di` is the incremental current flowing in each grid cell of the wire, and `dh` is the total magnetic field produced by that current at a distance of one meter.

From the figure, the difference in X coordinates between the cell in the wire and the cell in the area of interest is `dx` and the similar difference in Y coordinates is `dy`. Therefore, the distance between the two cells is

$$r = \sqrt{dx^2 + dy^2} \tag{A.3}$$

and the actual field generated in the cell in the area of interest is

$$H = \frac{dh}{r} \tag{A.4}$$

This field can be decomposed into two orthogonal vectors, one along the X axis (`xv`) and one

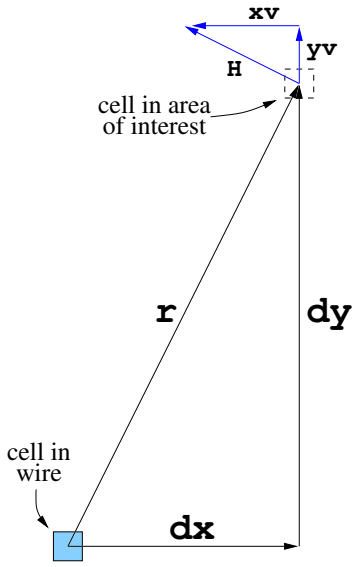

Figure A.2: Geometry of field calculations

along the Y axis (`yv`). The triangle formed by `xv`, `yv`, and `H` is a similar triangle to the one formed by `dy`, `dx`, and `r`, respectively, so

$$\frac{xv}{H} = \frac{dy}{r} \tag{A.5}$$

and

$$xv = \frac{dy}{r} \times H \tag{A.6}$$

$$= \frac{dy}{r} \times \frac{dh}{r} \tag{A.7}$$

$$= \frac{dy \cdot dh}{r^2} \tag{A.8}$$

$$= dy \cdot dh \cdot \left(\frac{1}{dx^2 + dy^2}\right) \tag{A.9}$$

By a similar analysis it can be shown that

$$yv = dx \cdot dh \cdot \left(\frac{1}{dx^2 + dy^2}\right) \tag{A.10}$$

In line 47 the X coordinate of the center of the first cell in the wire is computed, and line 48 then iterates over the X coordinates of all cells in the wire. Line 49 computes the Y coordinate of the first cell in the wire. Lines 50 and 51 calculate the value of dx and variable `dxsq` gets dx^2. In line 52 variable `dhdx` gets $dh \cdot dx$, which will be used to calculate `yv` later.

Line 53 begins a loop that iterates over all of the Y coordinates of the cells in the wire, and then dy is calculated in the next line. At this point the distance to the area of interest, r, can be calculated. However, the quantity that is actually needed is $\frac{1}{r^2}$, which is calculated as $\frac{1}{dx^2+dy^2}$ and saved in variable `irsq` in line 55. The field vectors produced by this cell in the wire are then calculated and accumulated as `xv` and `yv` in lines 56 and 57.

Line 58 moves the Y coordinate of the wire cell to the next cell and is the last line in the loop that started in line 53. Similarly, the X coordinate is advanced in line 59, the last line of the loop that started in line 48. At line 60 the program has iterated over all of the cells in the wire and calculated the net field in one cell of the area of interest, so the coordinates of that cell and its field component values are printed at this time.

Lines 61 through 64 simply keep track of the minimum and maximum values for the field components that are calculated over all cells in the area of interest. At line 65 the Y coordinate of the cell in the area of interest is advanced, ending the loop that started in line 44.

Line 66 adds a blank line to the data output, which is needed by the plotting software to separate horizontal lines of data points in the graph. Line 67 then advances the X coordinate of the cell in the area of interest and ends the loop that started in line 42.

Once the program exits the outermost loop and proceeds to line 69, the field values of all of the cells have been calculated and printed. Lines 69 and 70 print the maximum and minimum field values and the program is completed.

```
1 #!/usr/bin/python
2 """
3  fem
4
```

```
 5   Copyright 2005 by Kenneth Joseph Hass and the University of Idaho
 6
 7  """
 8  import os, math
 9  if __name__ == '__main__':
10    import sys
11  #
12  # current
13    i = 0.001
14  # wire dimensions
15    wminx = -0.5e-6
16    wmaxx =  0.5e-6
17    wminy = -2.0e-6
18    wmaxy = -0.0e-6
19  # area of interest
20    minx = -0.5e-6
21    maxx =  0.5e-6
22    miny =  0.002e-6
23    maxy =  0.003e-6
24  # grid size
25    step = 0.001e-6
26
27    di = i/(((wmaxx-wminx)/step)*((wmaxy-wminy)/step))
28    dh = di/(2.0*math.pi)
29    dh = dh*((4.0*math.pi)/1000.0)
30
31    print "# i=",i,"  di=",di,"  dh=",dh
32    print '# j= %10.4g A/m2' % (di/(step*step))
33    print "# wire from ",wminx,wminy,"  to ",wmaxx,wmaxy
34
35    xvmin = 1e6
36    xvmax = -1e6
37    yvmin = 1e6
38    yvmax = -1e6
39    halfstep = step/2.0
40
41    x = minx + halfstep
42    while x < maxx:
43      y = miny + halfstep
44      while y < maxy:
45        xv = 0.0
46        yv = 0.0
47        wx = wminx + halfstep
48        while wx < wmaxx:
49          wy = wminy + halfstep
50          dx = wx - x
51          dxsq = dx*dx
```

```
52          dhdx = dh*dx
53          while wy < wmaxy:
54            dy = y - wy
55            irsq = 1.0/(dxsq + dy*dy)
56            xv = xv + (dh*dy)*irsq
57            yv = yv + dhdx*irsq
58            wy = wy + step
59          wx = wx + step
60        print '%10.4g %10.4g %10.4g %10.4g' % (x, y, xv, yv)
61        if xv < xvmin: xvmin = xv
62        if xv > xvmax: xvmax = xv
63        if yv < yvmin: yvmin = yv
64        if yv > yvmax: yvmax = yv
65        y = y + step
66      print
67      x = x + step
68
69    print '# Maximum values      %10.4g %10.4g' % (xvmax, yvmax)
70    print '# Minimum values      %10.4g %10.4g' % (xvmin, yvmin)
```

Appendix B

MTJ SPICE Model

The simulation model used for a pair of MTJs is represented in the schematic of Fig. B.1 and the SPICE subcircuit listing below. The circuit has four terminals, which are the two terminals of each of the MTJs. To simplify the model, one of the MTJs is always assumed to be in the

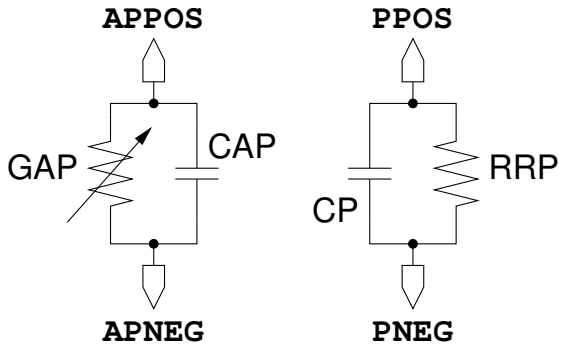

Figure B.1: MTJ SPICE model schematic

antiparallel state, with a higher resistance than the other MTJ, and no attempt is made to simulate the magnetic switching of the free layer. The bias dependence of the magnetoresistance ratio is modeled by varying the resistance of the antiparallel junction as a function of the voltage across it.

As shown in the figure, the parallel (low resistance) MTJ is modeled as a fixed resistance, `RRP`, in parallel with a fixed capacitance, `CP`. The antiparallel MTJ is modeled as an equal fixed capacitance, `CAP`, in parallel with a variable resistance, `GAP`.

Input parameters to the model include the junction area in square microns, `AREA`, the tunnel oxide thickness in meters, `TOX`, the resistance-area product of the junction in $\Omega\mu m^2$, `RA`, the magnetoresistance ratio at zero bias, `MRR`, the relative permittivity of the junction barrier, `ER`,

and a dimensionless smoothing parameter, SMOOTH. Each of these parameters has a default value, shown in line 3 of the subcircuit listing. The zero-bias resistances of the parallel and antiparallel junctions are calculated as RP and RPP in lines 11 and 12, and the parallel plate capacitance of the junction is calculated as CJ in line 13.

Lines 14 and 15 define the voltage-controlled resistance (VCR) of the antiparallel junction. In HSPICE a VCR is simply a special case of the 'G' element, a voltage-controlled current source. The voltage across the junction itself is used as the controlling parameter, with a piece-wise linear response that is symmetric about zero volts. When there is no voltage across the antiparallel junction it has the previously calculated zero-bias antiparallel resistance, which decreases smoothly to the zero-bias parallel resistance at an applied voltage of ± 0.5 V. Breakpoints equal to the zero-bias parallel resistance are added at ± 100 V to assure that the element's behavior is defined over its entire working voltage range. Lines 16 through 18 simply apply the previously calculated capacitance and resistance values to the other circuit elements.

```
1  * MTJ model with bias dependence
2  .subckt MTJ PPOS PNEG APPOS APNEG
3  + AREA=0.5 TOX=3.5n RA=5k MRR=1.3 ER=8.0 SMOOTH=0.05
4  *
5  * AREA is junction area in square microns
6  * TOX is oxide thickness in meters
7  * RA is resistance-area product in ohm-um^2
8  * MRR is magnetoresistance ratio at zero bias
9  * ER is the relative permittivity of the oxide, about 8 for Al2O3
10 *
11 .param RP='RA/AREA'
12 .param RAP='(RA/AREA)*MRR'
13 .param CJ='(8.854e-24*ER*AREA)/TOX'
14 GAP APPOS APNEG VCR PWL(1) APPOS APNEG
15 +   DELTA='SMOOTH' -100,'RP' -0.5,'RP' 0,'RAP' 0.5,'RP' 100,'RP'
16 CAP APPOS APNEG 'CJ'
17 RRP PPOS PNEG 'RP'
18 CP  PPOS PNEG 'CJ'
19 .ends
```

Appendix C

Pairs Plot Script for Gnuplot

Most of the graphs presented in this document were created using `gnuplot`, an open-source plotting program (see `www.gnuplot.info`). Many were straightforward, but the "pairs plot" graphs used some of the more esoteric capabilities of `gnuplot`. An example script to generate these plots is given here. In this example, a data file with five columns of numbers is to be plotted. The first four columns will be used as *X* and *Y* axis values for the data points, and the value in column five will determine whether a green, yellow, or red point is plotted.

A key element of the script is the idiom "`u 1:($5<L1?$2:1/0) ls 1`". The letter `u` is an abbreviation for `using`, which is a gnuplot keyword that specifies which columns of data will be used for the *X* and *Y* coordinates of the plot. The desired column numbers, separated by a colon (:), should follow the keyword. In this case the *X* coordinates will come from column one. The *Y* coordinates are obtained by evaluating the conditional expression within parentheses: "`$5 < L1 ? $2 : 1/0`". The meaning of this expression is "if the value in column five is less than the value of variable `L1` then use the value in column two as the *Y* coordinate, otherwise use the value of `1/0` as the *Y* coordinate". However, the expression `1/0` is invalid since division by zero is undefined. When `gnuplot` is given an invalid coordinate for a data point it silently ignores the data point, so the net result is that only those sets of *X* and *Y* coordinates for values in column five that are less than `L1` will be plotted. The last portion of the idiom, "`ls 1`", tells `gnuplot` to use "line style 1" to plot these points. This line style has been defined in the script to be green in color with filled circles to mark the data points.

The yellow and red data points are plotted in a similar fashion. For yellow points the conditional expression for the *Y* axis values is modified to only plot points between two

limits, stored in variables `L1` and `L2`, using line style two. Red points are plotted with a conditional expression that uses column five values greater than `L2` with line style three. The line styles were defined so that the green dots were larger than the yellow dots, which were larger than the red dots, but this is simply a personal preference. However, some graphics file formats (including postscript) do not support transparency so large dots of one color may obscure surrounding dots of another color.

An example of the datafile format is shown in Table C, with the resulting graph shown in Fig. C.1. Only the first few lines of data are shown; the graph displays 1000 data points.

Referring to listing below, lines 24 through 35 generate the three plots in the left column of the graph. The three plot commands in this section use the same column for the X coordinate and the next three columns for Y coordinates. Similarly, lines 36 through 43 create the two plots in the center column, using column two for the X coordinate versus columns three and four as the Y coordinate. Finally, column three is used for the X coordinate and column four as the Y coordinate in lines 44 through 47. Note that the size and placement of various elements, as well as the line style color definitions, are specifically for the encapsulated postscript output format and may need to be modified if another output format is used.

```
set terminal postscript eps enhanced color solid lw 1
set output "pairsplot.eps"
set multiplot
unset timestamp
unset key
unset xtics
unset ytics

L1 = 5.0
L2 = 10.0

set style line 1 lt 2  lw 1 pt 7 ps 0.3
set style line 2 lt 6  lw 1 pt 7 ps 0.2
set style line 3 lt 1  lw 1 pt 7 ps 0.1

set label "VARIABLE 1" at screen 0.10,0.01 center
set label "VARIABLE 2" at screen 0.24,0.21 center
set label "VARIABLE 3" at screen 0.38,0.41 center
set label "VARIABLE 2" at screen 0.02,0.13 center rotate
set label "VARIABLE 3" at screen 0.02,0.33 center rotate
set label "VARIABLE 4" at screen 0.02,0.53 center rotate

set size square 0.26,0.26
set origin 0.0,0.0
plot "file.dat" u 1:($5<L1?$2:1/0) ls 1, \
             '' u 1:(($5<L2&&$5>L1)?$2:1/0) ls 2, \
             '' u 1:($5>L2?$2:1/0) ls 3
set origin 0.0,0.2
plot "file.dat" u 1:($5<L1?$3:1/0) ls 1, \
             '' u 1:(($5<L2&&$5>L1)?$3:1/0) ls 2, \
             '' u 1:($5>L2?$3:1/0) ls 3
set origin 0.0,0.4
plot "file.dat" u 1:($5<L1?$4:1/0) ls 1, \
             '' u 1:(($5<L2&&$5>L1)?$4:1/0) ls 2, \
             '' u 1:($5>L2?$4:1/0) ls 3
set origin 0.14,0.2
plot "file.dat" u 2:($5<L1?$3:1/0) ls 1, \
             '' u 2:(($5<L2&&$5>L1)?$3:1/0) ls 2, \
             '' u 2:($5>L2?$3:1/0) ls 3
set origin 0.14,0.4
plot "file.dat" u 2:($5<L1?$4:1/0) ls 1, \
             '' u 2:(($5<L2&&$5>L1)?$4:1/0) ls 2, \
             '' u 2:($5>L2?$4:1/0) ls 3
set origin 0.28,0.4
plot "file.dat" u 3:($5<L1?$4:1/0) ls 1, \
             '' u 3:(($5<L2&&$5>L1)?$4:1/0) ls 2, \
             '' u 3:($5>L2?$4:1/0) ls 3
```

0.803	1.085	0.802	1.116	9.088
0.839	1.279	0.888	1.133	7.309
0.884	1.121	1.193	1.180	9.370
1.203	1.216	0.820	1.235	8.375
0.929	0.990	0.996	0.817	10.290
0.774	0.888	1.181	0.958	0.615
1.088	1.297	1.027	1.216	9.332
0.760	0.733	1.240	0.991	3.118
0.991	1.104	0.947	0.774	12.110

Table C.1: Pairs plot data example

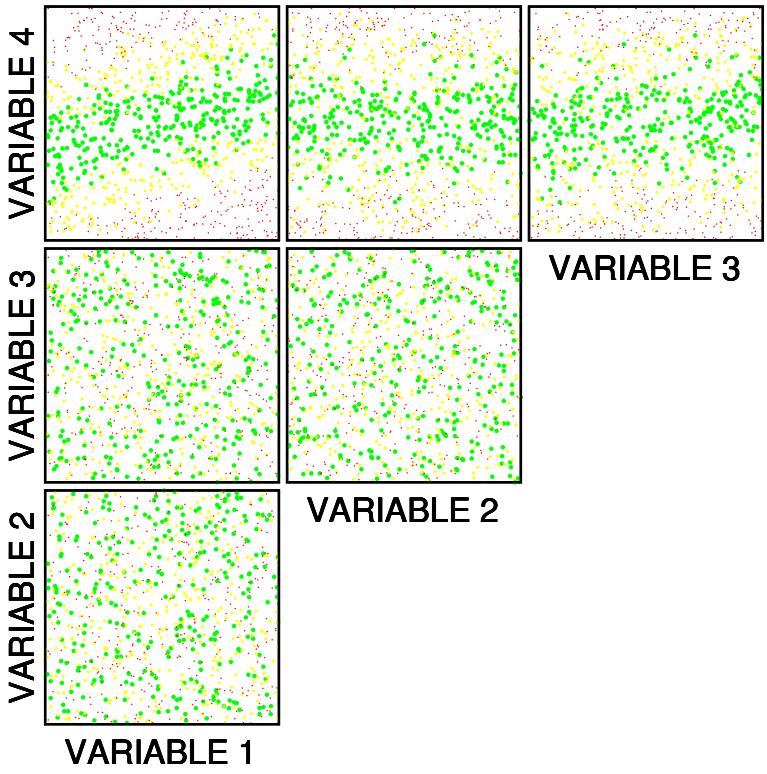

Figure C.1: Pairs plot graph example

Appendix D

First Shadow Latch Netlist

```
1  .subckt ShadowLatch1 Q QN NBG PBG VDD VSS
2  MQ0   VDU   LDN   VDD   PBG   pFLEXFET  L=0.18U  W=3.0U
3  MQ8P  Y3    Y2    VDU   VDU   pFLEXFET  L=0.18U  W=0.9U
4  MQ8N  Y3    Y0    VSS   BG13  nFLEXFET  L=0.18U  W=0.9U
5  MQ7P  Y2    Y1    VDU   VDU   pFLEXFET  L=0.18U  W=0.9U
6  MQ7N  Y2    Y3    VSS   BG02  nFLEXFET  L=0.18U  W=0.9U
7  MQ6P  Y1    Y0    VDU   VDU   pFLEXFET  L=0.18U  W=0.9U
8  MQ6N  Y1    Y2    VSS   BG13  nFLEXFET  L=0.18U  W=0.9U
9  MQ5P  Y0    Y3    VDU   VDU   pFLEXFET  L=0.18U  W=0.9U
10 MQ5N  Y0    Y1    VSS   BG02  nFLEXFET  L=0.18U  W=0.9U
11 MQ2P  Q     Y3    VDD   PBG   pFLEXFET  L=0.18U  W=2.55U
12 MQ2N  Q     Y3    VSS   NBG   nFLEXFET  L=0.18U  W=0.90U
13 MQ1P  Q     Y1    VDD   PBG   pFLEXFET  L=0.18U  W=2.55U
14 MQ1N  Q     Y1    VSS   NBG   nFLEXFET  L=0.18U  W=0.90U
15 MQ3P  QN    Y0    VDD   PBG   pFLEXFET  L=0.18U  W=2.55U
16 MQ3N  QN    Y0    VSS   NBG   nFLEXFET  L=0.18U  W=0.90U
17 MQ4P  QN    Y2    VDD   PBG   pFLEXFET  L=0.18U  W=2.55U
18 MQ4N  QN    Y2    VSS   NBG   nFLEXFET  L=0.18U  W=0.90U
19 .ends
```

Appendix E

Second Shadow Latch Netlist

```
1  .subckt ShadowLatch2 Q QN NBG PBG VDD VSS
2  MQ0   VDU   LDN   VDD   PBG   pFLEXFET  L=0.18U  W=3.0U
3  MQ8P  Y3    Y2    VDU   VDU   pFLEXFET  L=0.24U  W=0.9U
4  MQ8N  Y3    Y0    VSS   BG13  nFLEXFET  L=0.24U  W=1.47u
5  MQ7P  Y2    Y1    VDU   VDU   pFLEXFET  L=0.24U  W=0.9U
6  MQ7N  Y2    Y3    VSS   BG02  nFLEXFET  L=0.24U  W=1.47u
7  MQ6P  Y1    Y0    VDU   VDU   pFLEXFET  L=0.24U  W=0.9U
8  MQ6N  Y1    Y2    VSS   BG13  nFLEXFET  L=0.24U  W=1.47u
9  MQ5P  Y0    Y3    VDU   VDU   pFLEXFET  L=0.24U  W=0.9U
10 MQ5N  Y0    Y1    VSS   BG02  nFLEXFET  L=0.24U  W=1.47u
11 MQ2P  Q     Y3    VDD   PBG   pFLEXFET  L=0.18U  W=2.55U
12 MQ2N  Q     Y3    VSS   NBG   nFLEXFET  L=0.18U  W=0.90U
13 MQ1P  Q     Y1    VDD   PBG   pFLEXFET  L=0.18U  W=2.55U
14 MQ1N  Q     Y1    VSS   NBG   nFLEXFET  L=0.18U  W=0.90U
15 MQ3P  QN    Y0    VDD   PBG   pFLEXFET  L=0.18U  W=2.55U
16 MQ3N  QN    Y0    VSS   NBG   nFLEXFET  L=0.18U  W=0.90U
17 MQ4P  QN    Y2    VDD   PBG   pFLEXFET  L=0.18U  W=2.55U
18 MQ4N  QN    Y2    VSS   NBG   nFLEXFET  L=0.18U  W=0.90U
19 .ends
```

Appendix F

Third Shadow Latch Netlist

```
1  .subckt ShadowLatch3 Q QN NBG PBG VDD VSS
2  MQ0   VDU   LDN   VDD   PBG   pFLEXFET  L=0.18U  W=3.0U
3  MQ8P  Y3    Y2    VDU   VDU   pFLEXFET  L=0.24U  W=0.9U
4  MQ8N  Y3    Y0    VSS   BG13  nFLEXFET  L=0.24U  W=1.47u
5  MQ8C  BG13  Y0    BG13  BG13  nFLEXFET  L=0.24U  W=0.90u
6  MQ7P  Y2    Y1    VDU   VDU   pFLEXFET  L=0.24U  W=0.9U
7  MQ7N  Y2    Y3    VSS   BG02  nFLEXFET  L=0.24U  W=1.47u
8  MQ7C  BG02  Y3    BG02  BG02  nFLEXFET  L=0.24U  W=0.90u
9  MQ6P  Y1    Y0    VDU   VDU   pFLEXFET  L=0.24U  W=0.9U
10 MQ6N  Y1    Y2    VSS   BG13  nFLEXFET  L=0.24U  W=1.47u
11 MQ6C  BG13  Y2    BG13  BG13  nFLEXFET  L=0.24U  W=0.90u
12 MQ5P  Y0    Y3    VDU   VDU   pFLEXFET  L=0.24U  W=0.9U
13 MQ5N  Y0    Y1    VSS   BG02  nFLEXFET  L=0.24U  W=1.47u
14 MQ5C  BG02  Y1    BG02  BG02  nFLEXFET  L=0.24U  W=0.90u
15 MQ2P  Q     Y3    VDD   PBG   pFLEXFET  L=0.18U  W=2.55U
16 MQ2N  Q     Y3    VSS   NBG   nFLEXFET  L=0.18U  W=0.90U
17 MQ1P  Q     Y1    VDD   PBG   pFLEXFET  L=0.18U  W=2.55U
18 MQ1N  Q     Y1    VSS   NBG   nFLEXFET  L=0.18U  W=0.90U
19 MQ3P  QN    Y0    VDD   PBG   pFLEXFET  L=0.18U  W=2.55U
20 MQ3N  QN    Y0    VSS   NBG   nFLEXFET  L=0.18U  W=0.90U
21 MQ4P  QN    Y2    VDD   PBG   pFLEXFET  L=0.18U  W=2.55U
22 MQ4N  QN    Y2    VSS   NBG   nFLEXFET  L=0.18U  W=0.90U
23 .ends
```

DISTRIBUTION LIST

DTIC/OCP
8725 John J. Kingman Rd, Suite 0944
Ft Belvoir, VA 22060-6218 1 cy

AFRL/RVIL
Kirtland AFB, NM 87117-5776 2 cys

Official Record Copy
AFRL/RVSE/Jeffrey B. Scott 1 cy